

AD-A277 476



2

NAVAL POSTGRADUATE SCHOOL
Monterey, California



THESIS

DTIC
ELECTE
MAR 29 1994
S B D

IMPLEMENTATION AND EVALUATION OF AN
ASYNCHRONOUS GROUP MEMBERSHIP
PROTOCOL

by

David J. Pezdirtz, Jr.

December, 1993

Thesis Advisor:

Shridhar B. Shukla

Approved for public release; distribution is unlimited

94-09507



18612

DTIC 94-09507

94 3 28 059

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704	
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.</p>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE Dec 1993		3. REPORT TYPE AND DATES COVERED Master's Thesis, Final
4. TITLE AND SUBTITLE Implementation and Evaluation of an Asynchronous Group Membership Protocol			5. FUNDING NUMBERS	
6. AUTHOR(S) David J. Pezdirtz, Jr.				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE A	
<p>13. ABSTRACT (maximum 200 words) A group membership protocol provides the mechanisms to ensure the consistent group views among a group-oriented distributed processes. The protocol is required to dynamically re-configure the group views among the various members in the event of a change to the group due to a new member joining or a member departing. The departure may be voluntary or involuntary. The protocol must provide a scheme to detect the failure of any of the members and re-configure the group. Multiple changes to the group must be perceived at all members in the same order. This thesis deals with a particular group membership protocol. The protocol structures the group as a logical ring. Changes to the group are accomplished using a two-phase scheme. The agreement phase consists of circulation of an <i>agree</i> token. Processing the token makes a pending change known to all members. The commit phase incorporates the changes in the correct order. This thesis presents an implementation of this asynchronous group membership protocol. The main feature is that the decentralized nature of the protocol eliminates the need for a dedicated coordinator of changes. The processing requirements for the protocol are likewise distributed. The processing time required to implement a change is explored.</p>				
14. SUBJECT TERMS Asynchronous Group Membership Protocol, Unix based, distributed processing.			15. NUMBER OF PAGES 196	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Approved for public release; distribution is unlimited.

Implementation and Evaluation of an Asynchronous Group
Membership Protocol

by

David J. Pezdirtz, Jr.
Lieutenant, United States Navy
B.S.C.S., University of Vermont, 1983

Submitted in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE IN ELECTRICAL ENGINEERING

from the

NAVAL POSTGRADUATE SCHOOL

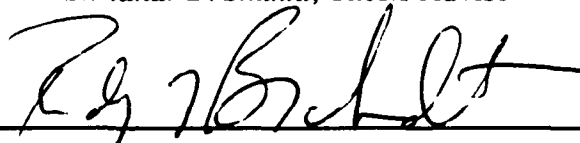
December 1993


Author:


David J. Pezdirtz, Jr.

Approved by:


Shridhar B. Shukla, Thesis Advisor


Randy L. Borchardt, Second Reader


Michael A. Morgan, Chairman
Department of Electrical and Computer Engineering

ABSTRACT

A group membership protocol provides the mechanisms to ensure the consistent group views among a group-oriented distributed processes. The protocol is required to dynamically re-configure the group views among the various members in the event of a change to the group due to a new member joining or a member departing. The departure may be voluntary or involuntary. The protocol must provide a scheme to detect the failure of any of the members and re-configure the group. Multiple changes to the group must be perceived at all members in the same order.

This thesis deals with a particular group membership protocol. The protocol structures the group as a logical ring. Changes to the group are accomplished using a two-phase scheme. The agreement phase consists of circulation of an *agree* token. Processing the token makes a pending change known to all members. The commit phase incorporates the changes in the correct order.

This thesis presents an implementation of this asynchronous group membership protocol. The main feature is that the decentralized nature of the protocol eliminates the need for a dedicated coordinator of changes. The processing requirements for the protocol are likewise distributed. The processing time required to implement a change to the group is shown to have a linear relationship to the group size.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
Re	
Distribution	
Availability	
Special	
Dist	
A-1	

TABLE OF CONTENTS

I. INTRODUCTION	1
A. BACKGROUND	1
1. Distributed Computing	1
2. Group Membership	1
B. SCOPE AND CONTRIBUTION	2
C. ORGANIZATION OF THE THESIS	2
II. GROUP MEMBERSHIP PROTOCOL	3
A. GROUP MEMBERSHIP PROTOCOL OVERVIEW	3
1. Assumptions	3
2. Overview	3
3. Processing of Individual Changes	5
a. Departure Processing	6
b. Join Processing	6
III. PROTOCOL CHANGES	7
A. LOST INITIAL PARAMETERS MESSAGE	7
1. Problem	7
2. Solution	7
3. Justification	8
4. Side effects	8
B. DUPLICATE PROCESSING	10
1. Problem	10
2. Solution	11
3. Justification	11
C. INVALID DELETE TOKEN	14
1. Problem	14
2. Solution	14
3. Justification	15
D. LOST TOKEN ACKNOWLEDGMENT	15
1. Problem	15
2. Solution	15
3. Justification	15
E. AGREEPROCESSOR	16
1. Initiate_Agreement Message	17

2. External Token Pool	18
a. Token Originator Failed	18
b. Duplicate Processing	20
3. Tokens	22
a. Joinreqst Tokens	22
b. Agree Tokens	22
c. Commit Tokens	23
4. Side Effects	23
F. MULTIPLE JOINS	24
1. Problem	24
2. Solution	25
3. Justification	26
G. OTHER IMPROVEMENTS	26
1. <i>Joinreqst</i> Token Processing	26
a. Problem	26
b. Solution	27
2. <i>Commit</i> Token Generation	27
a. Problem	27
b. Solution	27
3. Message Queue in the FIFO Channel Layer	28
H. SYNOPSIS	30
IV. PERFORMANCE OF THE GMP	31
A. LATENCY	31
B. TESTING	33
V. METHODS FOR IMPROVING PERFORMANCE	35
A. MESSAGE REDUCTION	35
1. TokenPool versus Tokens	35
2. Periodic Token Pool	36
3. Piggyback the Token Pool	36
B. SINGLE-THREADED PROGRAM	37
1. Problem	37
2. Solution	37
3. Justification	37
VI. CONCLUSIONS AND RECOMMENDATIONS	39

LIST OF REFERENCES	40
APPENDIX	41
INITIAL DISTRIBUTION LIST	185

LIST OF TABLES

Table 1: CONDITIONS TO DETECT DUPLICATE PROCESSING	21
Table 2: CONDITIONS WARRANTING A TIME STAMP	31
Table 3: PROCESSING TIME VALUES	33
Table 4: SINGLE THREADED PROCESSES AND EQUIVALENTS	38

LIST OF FIGURES

Figure 1: A Logical Ring	5
Figure 2: Reporting of Status	7
Figure 3: Integrate Member - Process Dependencies	9
Figure 4: Integrate Member Process Specification	9
Figure 5: Monitor Process - Internal Structure and Dependencies	10
Figure 6: Error Condition Arising from Asynchronous Programs	11
Figure 7: Agreement Processor - Process Dependencies	12
Figure 8: Commit Processor - Process Dependencies	13
Figure 9: Actions for Committing a Change	14
Figure 10: FIFO Channel - Back Process	16
Figure 11: Agreement Processor	17
Figure 12: Determination of Token Originator's Failure	19
Figure 13: Relative Rank	20
Figure 14: Processing Agree Tokens	23
Figure 15: External Token Pool Message Format	24
Figure 16: Simultaneous <i>Joins</i>	25
Figure 17: Group View after 1 member has joined	25
Figure 18: Group View at Host & J2	26
Figure 19: Processing of a Join Request Message / Token	27
Figure 20: Generate / Receive and Process a Commit Token	28
Figure 21: FIFO Channel - Front Process	29
Figure 22: Group Changes Required by the Filter Program	32

Figure 23: Time Stamp File Format	32
Figure 24: Average Time for each Member to Implement a Join to the Group	34
Figure 25: Average Time to Implement a Failure in the Group	34
Figure 26: Single-Threaded Process Inter-Dependencies	37

ACKNOWLEDGEMENT

Even though this thesis bears my name, there are two people who contributed considerably to its completion . In addition to my efforts, the original code was written and developed by Prof. Shridhar Shukla and LT Fernando Pires, Portuguese Navy. Their help and cooperation made this project significantly smoother.

Finally, I would like thank my wife Sandi. Without her support and tolerance this thesis would not have been possible. To her, I dedicate this thesis.

I. INTRODUCTION

A. BACKGROUND

1. Distributed Computing

Distributed computing is at the forefront of today's computing research. The increased reliability and performance has resulted in distributed computing being used in various applications such as distributed control applications, distributed databases, and real-time settings [1]. The need for fault-tolerant systems is particularly important to military applications. Such applications typically require fault tolerant algorithms, real-time response, on-line reconfiguration, and other schemes to increase reliability. These requirements, however, lead to significant additional complexity. This complexity arises in part due to the reliable inter-processor communication required to implement the distributed processing. Networks are inherently unreliable and make reliable application level message passing a non-trivial task. One of the primary requirements of reliable distributed applications is a reliable multicast communication primitive. A group membership protocol simplifies the construction of a such a primitive [2].

2. Group Membership

Cooperating processes constituting a single application share resources and constitute a process group. Underlying the consistent behavior of such a group is the requirement that all members of the group have implicit knowledge of all other members in the group. Additionally, all members must perceive changes to the group in the same order. Group membership will change as processes join or depart the group. Departures may be voluntary or involuntary as in the case of a failure. An additional requirement is the timely detection of members failing. In order to enhance the robustness of the system, intra-member monitoring must occur. This can be a simple exchange of messages indicating that the process being monitored is still "live."

Historically, protocols solving the group membership problem have been of a centralized design. One member acted as the group host and had the responsibility of monitoring all subordinate members. All changes were detected by the host. Upon a change to the group view, the host broadcast the new group view to all members of the group. Obviously, problems can arise if the host itself fails. Voting must occur in order to elect a new host with the added overhead associated with the voting. Additionally, the processing requirements are unequally distributed between the members of the group as the host takes a major share of the load.

B. SCOPE AND CONTRIBUTION

In this thesis, an implementation of the decentralized asynchronous membership protocol is presented. The protocol was originally presented in [3] and [4]. It was further refined and partially implemented in [5]. This thesis presents a brief overview to the protocol. A more detailed explanation of the protocol and definitions can be found in [5]. This thesis covers additional refinements to the protocol necessary to correct flaws discovered in the coding and testing phase of development. Additionally, the performance of the protocol on an Ethernet Local Area Network (LAN) is presented and discussed.

C. ORGANIZATION OF THE THESIS

The thesis is divided into six chapters. Chapter II presents an overview of the basic operation of the protocol. Chapter III discusses the changes required to implement improvements to the protocol. Chapter IV covers possible further refinements to the protocol. In Chapter V, the performance of the current protocol is analyzed. Chapter VI presents the analysis and possible future areas of research. The code developed is included in the Appendix.

II. GROUP MEMBERSHIP PROTOCOL

In this chapter, the group membership protocol (GMP) is described. The original protocol was presented in [3] and further developed in [5]. This chapter presents an overview of the protocol.

A. GROUP MEMBERSHIP PROTOCOL OVERVIEW

1. Assumptions

The following assumptions are made by the GMP in order for proper implementation. A fully-connected network of reliable First-In / First-Out (FIFO) communication channels connecting operational members is assumed. All failures are assumed to be *crash* or *fail-stop*[6]. This implies that a message sent will be delivered unless the recipient has failed. However, there is no upper bound on the time of delivery.

Multiple changes to the membership are allowed simultaneously. However, the changes are committed one at a time and in the same order at all members.

A member is added to the group when a *join* is processed and is removed when a *failure* is perceived.

The group name is assumed to be public. Those elements that may wish to become members by joining the group have access to the group file which contains the current group view. A prospective member searches for the file on a given site. The group view is extracted and the prospective member sends the *joinreqst* to the appropriate address.

The protocol maintains three main databases at each member; the membership list (group view), status table and token pool. Each has a separate database manager to ensure mutual exclusion to all processes needing the services of the database.

2. Overview

The GMP guarantees that the group view changes occur in the same relative sequence at all operational members of the group.

The most significant feature of the GMP is the decentralization. No single member is responsible for detecting a change to the membership nor guaranteeing group view consistency among the group members. A logical ring is used to implement both of these functions in a distributed manner. The logical ring is a circular ordering of the members of the group.

The physical location of the member has no relation to the ordering of the logical ring. Within the ring structure the direction of traversal was arbitrarily chosen as clockwise. Given the structure each member only monitors its anti-clockwise neighbor, the *acwnbr*. The *acwnbr* responds to a status query, *statusqry*, with a status report, *statusrpt*. Likewise, it sends a *statusqry* to its *acwnbr*. Thus, every member monitors only one other member of the group and is itself monitored by its clockwise neighbor, *cwnbr*.

Consider a five member group. Process p_0 was the initial member of the group. The other members joined in an order such that p_0 is the *acwnbr* of p_1 , p_1 is the *acwnbr* of p_2 , and so on. Member p_1 sends a *statusqry* to p_0 , which responds with a *statusrpt*. Likewise, p_2 sends a *statusqry* to p_1 . This is illustrated in Figure 1. For clarity, only the monitoring and response is shown for the first set of neighbors.

The ring configuration changes as the group membership changes. The ring starts as a single member group with other members joining in some arbitrary order. All changes to the group are treated in a similar manner. Members wishing to join an existing group read the group membership file in the first active member located on the net. The joining member then sends a *joinreqst* to the first member of the group. If the initial parameters message, *initparams*, is not received in a reasonable time, the request is transmitted to the next member in the group view file until all members have been attempted or a successful join has been completed.

A failure is considered an involuntary departure. When a member departs the group voluntarily, it simply stops responding to *statusqrys*. It will then be perceived as failed and subsequently removed from the group. Delayed transmission of the *statusrpt*,

failure of the member to respond to the *statusqry* and a lost *statusrpt* will all result in the monitored member being detected as failed.




Figure 1 A Logical Ring

3. Processing of Individual Changes

The GMP allows for a two phase procedure for all changes to the group view, the *agree* phase and the *commit* phase. An *agree* token is circulated around the ring. Once the originator receives the token via the ring, all members have processed the *agree* token. At this point the *agree* token is converted to a *commit* token, and the change agreed upon is committed by each member as the token is circulated around the ring. The protocol ensures that each token is received by all members, processed only once, and never lost. More complete descriptions of the actions required by the different phases is covered in the following section.

a. Departure Processing

Once a member perceives the departure of its *acwnbr*, voluntary or otherwise, a *failagree* token is initiated. The failed member is added to the status table with a *failagree* status. The token is incorporated into the token pool and transmitted via the FIFO channel to the *cwnbr*. Similar processing occurs at members receiving a *failagree* token for the first time. Once the token has been received by the originator, the agree phase has been completed. The failed member is removed from the status table and group view. The token pool is purged. The *failcomit* token added to the token pool and transmitted around the ring. When a member receives a *failcomit* token for the first time, similar processing occurs.

b. Join Processing

The protocol maintains a logical marker between the first and last members to join the group. The first member is called the *host*. A new member will always join the group as the *acwnbr* of the *host*. The *host* has the responsibility of initiating the *joinagree* token for the new member. However, the *host* may not be the member of the group that receives a join request message from the new member. In this case, the message is converted to a token, forwarded, and the new member is added to the status table. Once the *host* has received either a *joinreqst* token or message, it initiates a *joinagree* token. The *host* then adds the token to the token pool, the member to the status table, and transmits the token to its *cwnbr*. Similar processing occurs when a member receives a *joinagree* token for the first time. When the *host* receives the *joinagree* token via the token ring, it initiates the commit phase.

The *joincomit* phase consists of purging the token pool, incorporating the *joincomit* token in the token pool, adding the new member to the group view, and forwarding the token. Additionally, the host will transmit the status table, group view, and token pool to the new member in the initial parameters message, *initparams*. This is accomplished by *IntegrateMember*.

III. PROTOCOL CHANGES

This chapter describes the revisions to the group membership protocol proposed in [5]. Modifications discussed include lost messages, delayed transmission, lost token acknowledgments, and proper termination of the agree phase.

A. LOST INITIAL PARAMETERS MESSAGE

1. Problem

Consider a lost initial parameters *initparam* message. After the *initparam* message is sent, the joining member is regarded as part of the group by the sender even if it is lost. However, the group membership, token pool and status table are not accurately reflected in the joining member's local database. There was no mechanism for re-transmission of the *initparam* message, nor was it possible to recreate the information locally.

2. Solution

The status reporter is required authenticate the group membership. Reports are generated in response to status queries from only those members that are in the group or are joining the group. Status queries from outside the group are ignored. See Figure 2.

ReportStatus process at p_i

```
1  if (not blocked by IntegrateMember)
2    if (querying member  $\in GV_{p_i}$  or has joinagree status)
3       $p_{mon} =$  querying member
4      send status to  $p_{mon}$ 
5      if (previous querying member =  $p_{mon}$ )
6        send TokenPool( $p_i$ ) to  $p_{mon}$ 
7      end
8    end
9  end
end ReportStatus
```

Figure 2 Reporting of Status

3. Justification

A lost *initparam* message will result in the new member failing to respond to the host's first *statusqry*. Upon time out, the new member will be considered a failed process. The host's original *acwnbr* will then be monitored anew by the host. The new member, never having received the *initparam* message will time out on the join request. It will then attempt to join again.

By the time the *initparam* message is transmitted, the host's original *acwnbr* has knowledge that the new member is now joining the group. Therefore, with this change, the host's original *acwnbr* issues status reports to the new member's queries.

4. Side effects

Such group authentication will prevent multiple switching of the *cwnbr*. If the *initparam* message is lost, the host's original *acwnbr* will be un-monitored for a short time. Upon the failure detection of the new member, the host's original *acwnbr* will again be monitored by the host process. To ensure that the *StatusReporter* responds only to members within the current group view, *IntegrateMember* must be atomic with respect to *StatusReporter*. This will prevent race conditions when *initparam* message is received, followed by an almost simultaneous receipt of the first status query from the host process.

Figure 3 shows the process dependencies, while Figure 4 depicts the specification for *IntegrateMember*. The inter-process dependencies for the monitor processes are shown in Figure 5.




Figure 3 Integrate Member - Process Dependencies

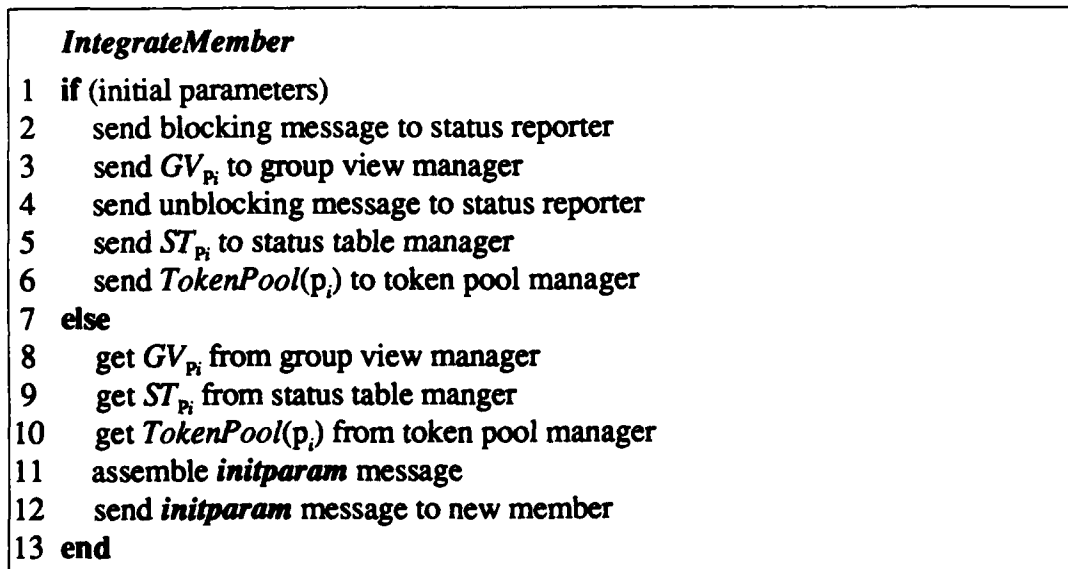


Figure 4 Integrate Member Process Specification

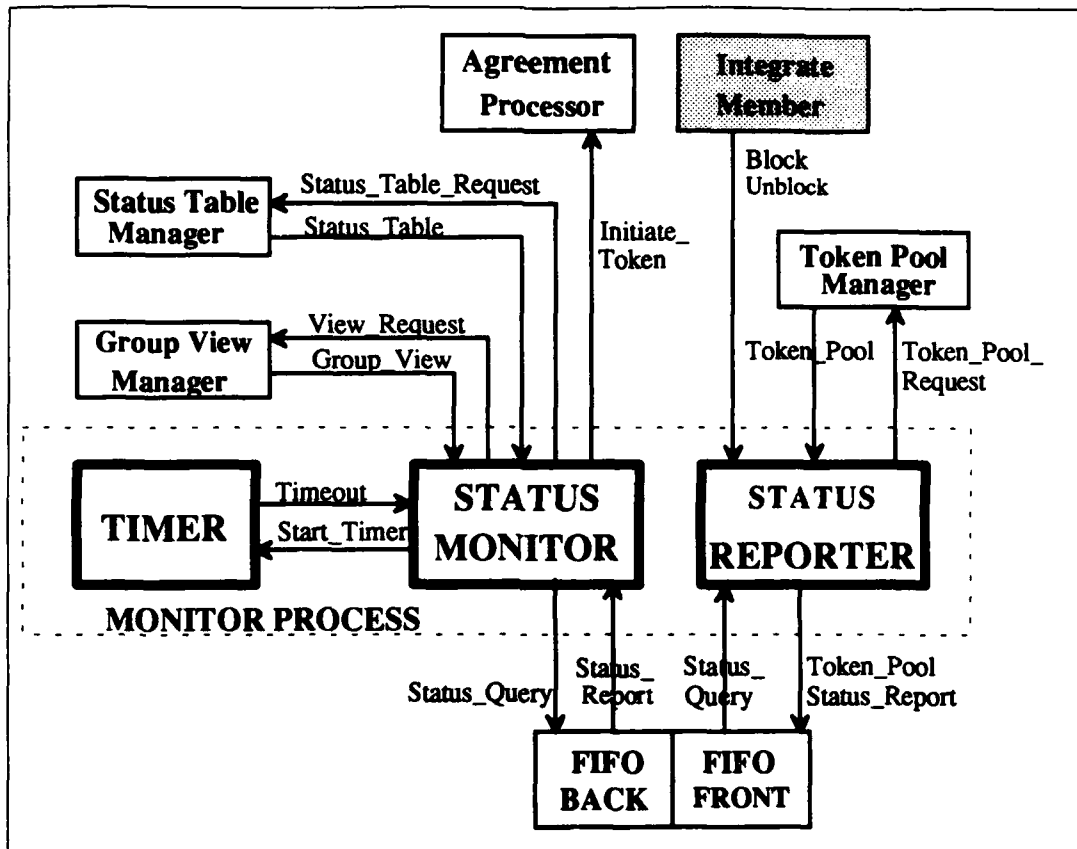


Figure 5 Monitor Process - Internal Structure and Dependencies

B. DUPLICATE PROCESSING

1. Problem

The protocol was designed for processing in a member to be concurrent. Consider *AgreeProcessor* and *ComitProcessor*. It is possible for the *AgreeProcessor* to receive an *agree* token immediately followed by an external token pool containing the same token. The token may be converted to a *commit* token and forwarded to *ComitProcessor*. Due to context switching, processing of the *commit* token may not be immediate. Further, *AgreeProcessor* begins processing the token pool which contains the copy of the original *agree* token. Since the processing of the *commit* token has not occurred, it is possible for the *agree* token from the token pool to be detected as requiring

conversion. The subsequent processing of the duplicate commit is an error. Figure 6 depicts the problem.

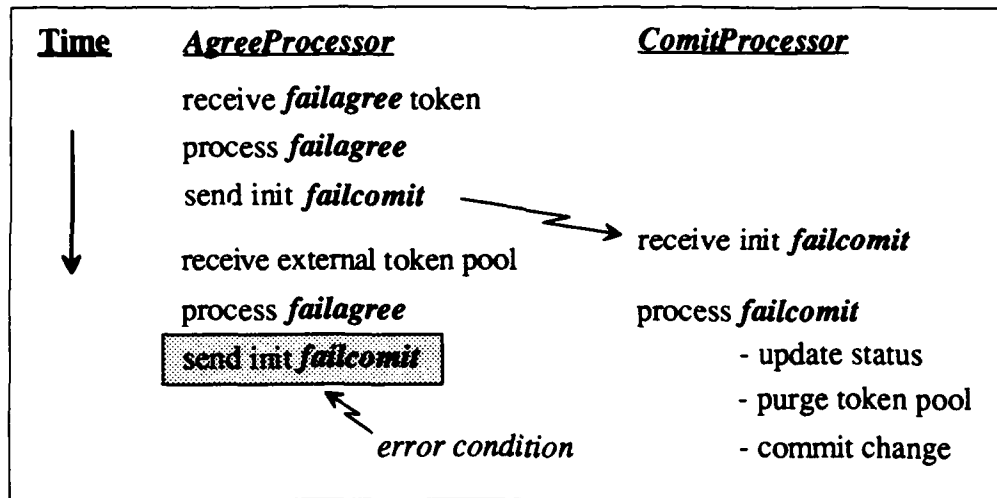


Figure 6 Error Condition Arising from Asynchronous Programs

2. Solution

In order to solve this problem, *ComitProcessor* must be atomic with respect to *AgreeProcessor*. (see Figures 7 and 8)

3. Justification

Both processors use the same databases. Rejection of duplicate tokens depends upon the current state of the databases. Since token rejection is accomplished by *AgreeProcessor*, and *ComitProcessor* updates the state to reflect the *commit* in progress, *AgreeProcessor* must not begin its next iteration until *ComitProcessor* has updated the state. *ComitProcessor* does not fully update the state until just prior to transmitting the *commit* token around the ring. Refer to Figure 9, lines 1-3.

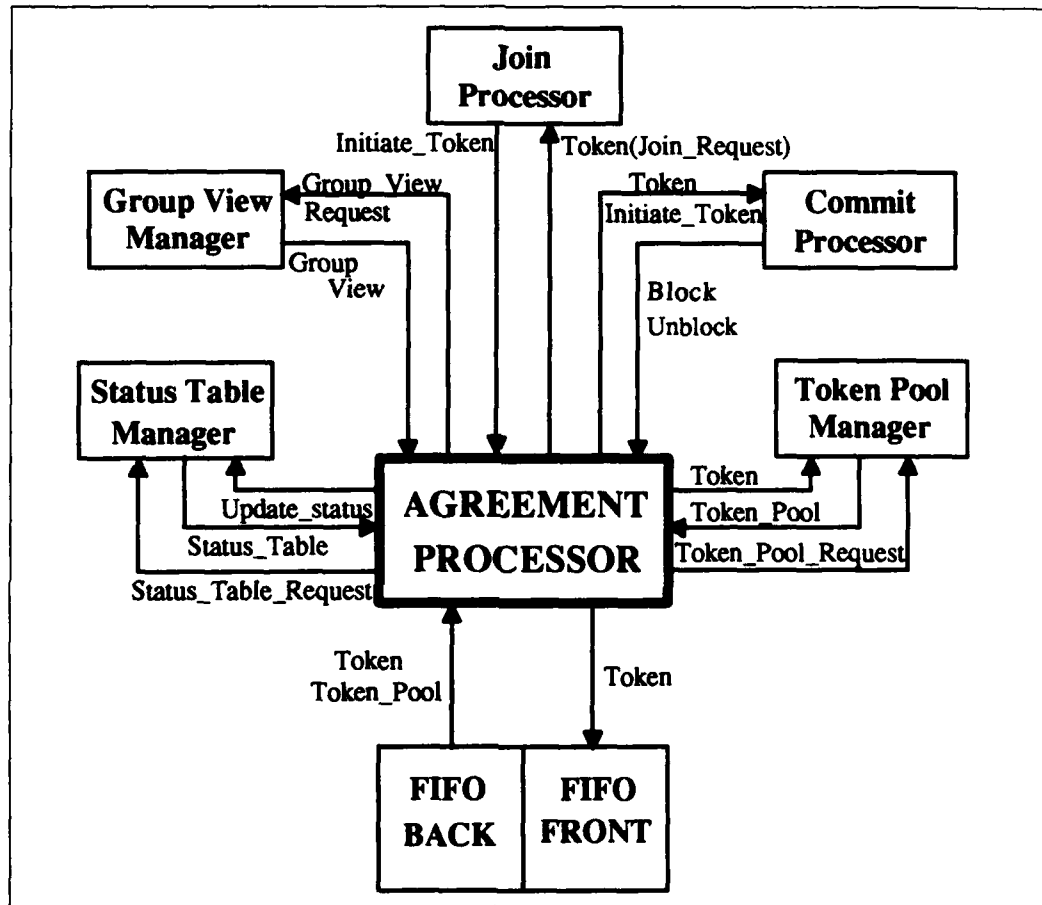


Figure 7 Agreement Processor - Process Dependencies

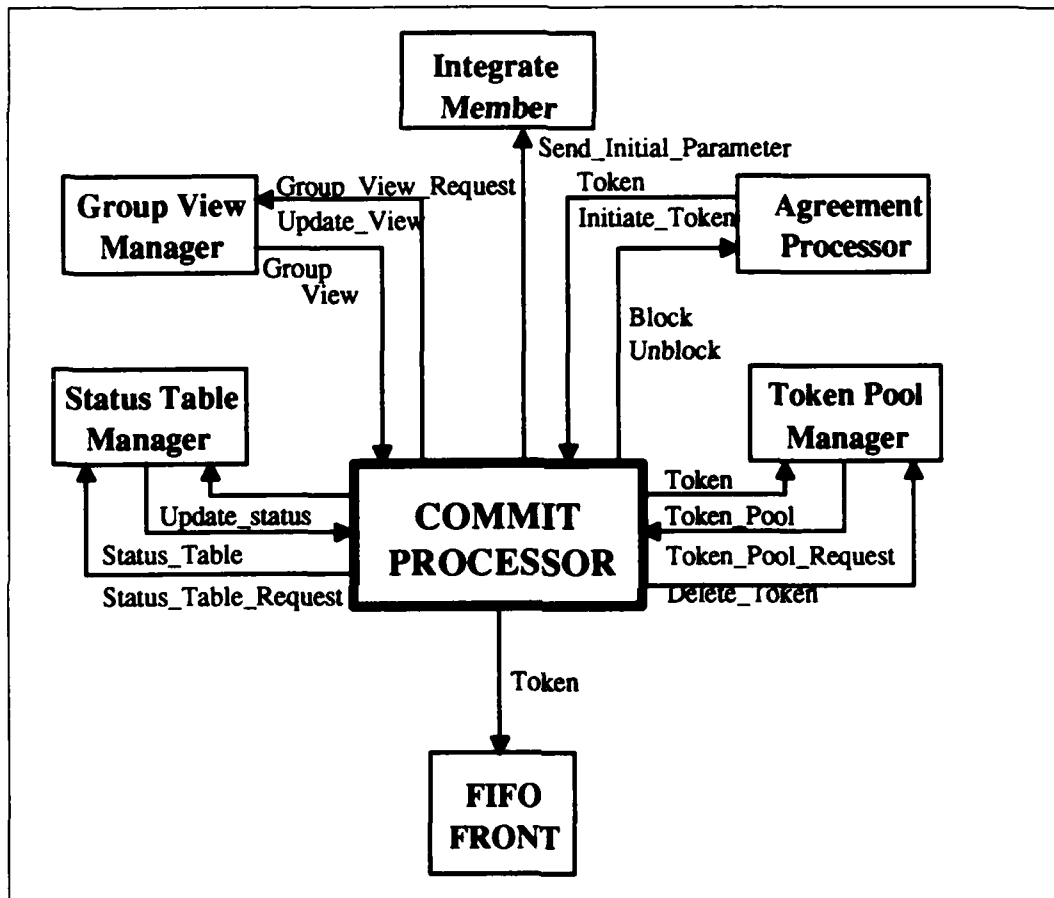


Figure 8 Commit Processor - Process Dependencies


```

CommitChange for  $commit_{p_j}(p_k)$  at  $p_i$ 
/* Depending on whether a join or departure */
1  add or delete  $p_k$  from  $GV_{p_i}$ 
2  delete  $p_k$  entry from  $ST_{p_i}$ 
3   $vn(p_i) \leftarrow vn(p_i) + 1$ 
4  delete all commit tokens received before  $agree_{p_j}(p_k)$  from  $TokenPool(p_i)$ 
5  if (join committed &&  $joinreq_{p_j}(p_k) \in TokenPool(p_i)$  )
6    delete  $joinreq_{p_j}(p_k)$ 
7  end
8  add  $commit_{p_j}(p_k)$  to  $TokenPool(p_i)$ 
9  delete  $agree_{p_j}(p_k)$ 
10 if (current host =  $p_k$ )
11   determine new  $p_{host}$ 
12 end
13 if ((join committed) && ( $p_{host} = p_i$ ))
14   send  $ST_{p_i}$ ,  $TokenPool(p_i)$ , and  $GV_{p_i}$  to  $cwnbr(p_i)$ 
15 end
16 send  $commit_{p_j}(p_k)$  token to  $cwnbr(p_i)$ 

end CommitChange

```

Figure 9 Actions for Committing a Change

C. INVALID DELETE TOKEN

1. Problem

An attempt to delete a nonexistent *joinreqst* token from the token pool would result in an error condition and would hang the process. The *joinreqst* token is not always present in the token pool. The token occurs only if a joining member has made the request to a member that is not the host.

2. Solution

Before attempting to delete the *joinreqst* token, the token pool is checked. If the token is present, it is then deleted. Figure 9, lines 5-7.

3. Justification

This is a special case token, and does not always occur in all token pools. Exception handling as above will correct any inconsistencies among the various token pools.

D. LOST TOKEN ACKNOWLEDGMENT

1. Problem

If the token acknowledgment is not received by the sending *front* process, the token will remain on the queue. It will be re-transmitted every time that *front* receives a message. The original token will be processed at the receiving end and further receipts of the same token will be detected as duplicates and discarded. The problem lies in that the queue is never cleared. Therefore, subsequent tokens will be blocked behind the token for which the acknowledgment was lost, unless a failure of one of the two members occurs.

2. Solution

This problem is solved by checking the serial number of the message on the receiving end of the FIFO channel. If the message is the last token received, the token acknowledgment is re-transmitted.

3. Justification

If the message received is the expected token, an acknowledgment is sent. The token is forwarded to the appropriate internal sub-process. If the last token received is received again, a token acknowledgment is sent back and the duplicate token is discarded. This mechanism will account and correct for lost token acknowledgments. Figure 10, lines 17-19.

FIFO Channel - *BACK* process

```
1  Wait for a channel ready to ready
2  if (internal channel ready)
3      if (Status_Query)
4          update acwnbr
5          send Status_Query
6  else if (Initial_Parameters)
7          update acwnbr
8          send Initial_Parameters
9  else if (Join_Request)
10     send Join_Request
11  end
12 else /* external channel ready */
13     if (message originator = acwnbr)
14         if (Status_Report)
15             send Status_Report to MONITOR_PROCESS
16         else if (Token)
17             if (Serial_Number = Expected_Serial_Number - 1)
18                 send Token_Ack /* to acwnbr */
19             end
20             if (Serial_Number = Expected_Serial_Number )
21                 send Token to AgreeProcessor
22                 send Token_Ack /* to acwnbr */
23                 increment Expected_Serial_Number
24             end /* out of order messages are discarded */
25         else if (Token_Pool) /* Token_Pool is always accepted */
26             send Token_Pool to AgreeProcessor
27             send Token_Ack /* to acwnbr */
28             set Expected_Serial_Number = Serial_Number + 1
29         end
30     end
31 end
```

Figure 10 FIFO Channel - Back Process

E. AGREEPROCESSOR

The specification for the *AgreeProcessor* was rewritten to account for various subtleties and to improve overall readability. All tokens received through the FIFO channel layer are sent to *AgreeProcessor* for dispatching to the appropriate processor. A

duplicate token is one that has been previously processed at given member. Some scheme to detect and reject duplicate tokens is required. Proper termination of the *agree* phase and subsequent initiation of the *commit* phase are also required. Figure 11. In this section, we discuss the operation of the *agree* processor.

```

AgreeProcessor for  $agree_{p_i}(p_k)$  at  $p_i$ 
1  if (not blocked by CommitProcessor)
2    if (initiate agreement message received) /*  $p_i = p_j$  */
3      add  $agree_{p_i}(p_k)$  to  $TokenPool(p_i)$ 
4       $ST_{p_i}(p_k) \leftarrow joinagreed$  or  $failagreed$ 
5      send  $agree_{p_i}(p_k)$  to  $cwnbr(p_i)$ 
6      send acknowledgment to calling process
7    else /* a token or external token pool is received */
8      if ( $ExtTokenPool$ )
9        for  $\forall tokens \in ExtTokenPool$ 
10         if ( $token \in TokenPool(p_i)$ )
11           if (originator failed)
12             ProcessToken
13           end
14         else /* token not in  $TokenPool$  */
15           if (received for the first time)
16             ProcessToken
17           end
18         end
19       end
20     else /* a token was received */
21       if (received for the first time)
22         ProcessToken
23       end
24     end
25   end
26 end

```

Figure 11 Agreement Processor

1. Initiate_Agreement Message

When *AgreeProcessor* receives an *initiate_agreement* message, the appropriate *agree* token is generated, added to the token pool and forwarded to the *cwnbr*. The

status table entry for the subject is updated. An acknowledgment is returned to the calling process. This reflects no change to the prior specification.

2. External Token Pool

When an external token pool is received, it is compared to the local token pool. All tokens in the external token pool are examined. Processing of a token depends on whether the token is also present in the local token pool.

If an *agree* token from the external token pool is in the local token pool, the token originator may have failed. Due to a failure of the originator, the *agree* token is requires conversion into a *commit* token at the first active clockwise neighbor of the originator only. Such tokens are processed as if they had been received as a separate token message. If the token is not present, it may have already been purged. These tokens must be rejected as duplicates.

a. Token Originator Failed

Detection of the token originator failing prior to initiating the *commit* phase is accomplished separately for *joinagree* and *failagree* tokens. Figure 12. It is necessary for the next active member to detect the originator's failure and also initiate the *commit* phase for the incomplete change started by the originator. The *agree* token may be received as part of the external token pool. The token will also be present in the local token pool from the *acwnbr* of the failed originator. This situation may also occur if a member in the middle of the ring fails. The failed member's *cwnbr* will receive an external token pool containing the original *agree*. However, this does not require a *commit* to be initiated as the originator has not failed. It is essential that the differences be noted and accounted for. The same conditions are present, but different processing must occur.

The originator's failure during a *joinagree* phase can be detected if the rank of the external token pool originator is greater than the rank of the current member. Consider the host failing prior to initiating the *commit* phase. All members in the group have agreed to the *join*. The *joinagree* token will be received by the new host upon ring reconfiguration. The new host's rank is zero (0) while the external token pool originator's

rank is the (group size - 1). Since $\text{Rank}(\text{originator}) > \text{Rank}(\text{host})$, a *commit* must be initiated.

Consider member p_i with rank i failing during the *joinagree* phase. When the failure of p_i is detected, p_{i+1} receives the external token pool from p_{i-1} . The *joinagree* token is present in both the external token pool and $\text{TokenPool}(p_{i+1})$. The token is rejected since $\text{Rank}(p_{i-1}) < \text{Rank}(p_{i+1})$.

```

LostAgreeToken
1  if (joinagree)
2    if ( $\text{rank}(p_j) > \text{rank}(p_i)$ )
3      return true
4    else
5      return false
6    end
7  end

8  if (failagree)
9    if ( $\text{RelativeRank}(p_k, p_i) > \text{RelativeRank}(p_j, p_i)$ )
10     return true
11   else
12     return false
13   end
14 end

```

Figure 12 Determination of Token Originator's Failure

Now we consider a duplicate *failagree* token. Define $\text{RelativeRank}(p_j, p_i)$ as the rank of p_j with respect to p_i instead of the host. A ring transversal starts and ends at the same specified process, i.e. any given member follows itself in a ring transversal. Figure 13. RelativeRank for a process that is not a member of the group is undefined. Recall the subject of a *failagree* remains a member of the group view until the *commit* is processed. A lost *failagree* token is determined by the RelativeRank of the token subject p_j and token pool originator p_{p_0} .




Figure 13 Relative Rank

If the $RelativeRank(p_i, p_i) \geq RelativeRank(p_{po}, p_i)$ the token originator has failed and the *failcommit* phase should be initiated. This situation will occur if the failed token originator itself fails prior to initiating the *commit* phase.

b. Duplicate Processing

Conditions to detect if a token has been received and processed already are summarized in Table 1. It is necessary to detect duplicate processing if the external token pool contains tokens not found in the local token pool. There are two possible ways for this to occur. If the token has not been received and processed, or the token has been purged from the local token pool.

Table 1 CONDITIONS TO DETECT DUPLICATE PROCESSING

<u>Token</u>	<u>Condition</u>
<i>joinreqst</i> <i>joinagree</i> <i>joincomit</i>	$p_k \in GV_{p_i}$
<i>failagree</i> <i>failcomit</i>	$p_k \notin GV_{p_i}$

Consider the *commit* phase. All members of the group have agreed to a particular ring reconfiguration. Due to the latency of token transmission around the ring, not all members commit the change simultaneously. Recall the token pool is purged of old *commit* tokens and the corresponding *agree* token prior to the *commit* token transmission. If a member that has committed a change receives an external token pool from a member that has agreed to the change, the *agree* token remains in the external token pool. Likewise, a previous *commit* token may be received as part of the external token pool. Since the tokens have been processed and removed from the local token pool, it is necessary to check the effects these tokens may have had on the group view, had they been previously processed. For duplicate *join* tokens, the subject would already be a part of the group view. Conversely, the subject of duplicate *fail* tokens would have already been removed from the group view. In this manner, duplicate processing can be detected and the tokens rejected.

It is not necessary to include all possible conditions for a token having been processed. Recall the duplicate processing check occurs only if the token is not present in the local token pool. For a *joinreqst*, if the token is received as part of the external token pool and has been purged from the local token pool, the *joincomit* must have occurred. Since the result of the *joincomit* is subject becoming a part of the group, it is only necessary to check the end result. Intermediate stages of the *join* process have not deleted the *joinreqst* token from the local token pool. Since the token remains in the

token pool, it is in both local and external token pools and is rejected. Similar logic results in the conditions presented in Table 1.

3. Tokens

Each type of token is handled individually upon receipt by *AgreeProcessor*. Figure 14. *AgreeProcessor* acts as a filter to remove duplicate tokens before forwarding non-*agree* tokens to the appropriate processor. *Agree* tokens are processed locally.

a. *Joinreqst Tokens*

The *joinreqst* tokens are forwarded to *JoinProcessor* for further processing.

b. *Agree Tokens*

If the current process is not the originator of the token, and it is not part of the local token pool, it is added to the token pool, the status updated and the token sent to the *cwnbr*. This accounts for the first time an *agree* token is received and processed.

```

ProcessToken
1  if (joinreqst)
2      send token to JoinProcessor
3  elseif (commit)
4      send token to ComitProcessor
5  elseif (agree)
6      if ( $(p_i \neq p_j) \ \&\& \ (\text{agree token} \notin \text{TokenPool}(p_i))$ )
7          add  $\text{agree}_{p_j}(p_k)$  to  $\text{TokenPool}(p_i)$ 
8           $ST_{p_i}(p_k) \leftarrow \text{FailAgreed or JoinAgreed}$ 
9          send  $\text{agree}_{p_j}(p_k)$  to  $\text{cwnbr}(p_i)$ 
10     else
11          $p_j$ 
12         if ( $((p_i = p_j) \parallel (\forall p_l \mid p_l \rightarrow p_i, p_l \in ST_{p_i}))$ )
13             compute rank  $\forall p_l \in ST_{p_i}$  with Agreed status
14             if rank( $p_i$ ) = smallest
15                 send initiate_comit to ComitProcessor
16             else
17                  $ST_{p_i}(p_k) \leftarrow \text{joinpendg or failpendg}$ 
18             end
19         end
20     end

```

Figure 14 Processing Agree Tokens

If the current process is the token originator, or the first active process clockwise from the originator that receives the *agree* token after it circulates around the ring, the rank of all processes with an *agreed* status is computed. An *initiate_commit* is transmitted if the subject is the lowest ranked of all processes fulfilling the above conditions. If not, the subject's status is updated to *pending*.

c. Commit Tokens

A *commit* token is immediately sent to *CommitProcessor* for processing.

4. Side Effects

The external token pool mentioned above requires a new message type, *ExtknPool*. Figure 15. The message includes the originator of the token pool as an internal field. This is required by *AgreeProcessor* for determination of a failed token originator.




Figure 15 External Token Pool Message Format

F. MULTIPLE JOINS

1. Problem

Consider two members attempting to join a group almost simultaneously. Figures 16-18. The first member's *join* will be completed properly. Upon completion of the first *join*, it is possible to begin processing the second member's *joinreqst* before the FIFO channels reflect the re-configured ring with the first member in it. It is possible to complete the second *join* before the channels are re-configured. This can happen because of the de-coupled protocol and FIFO channels. The host does not change its *acwnbr* until it initiates a status query to the last member. The FIFO channel determines a member's *acwnbr* as the target of the most recent *statusqry*. The *cwnbr* is the originator of the most recent *statusqry* received. There is an inherent latency involved in the FIFO channel reconfiguration due to the timing considerations of subsequent *statusqrys*. Thus, it is possible for a *joinreqst* from a second new member to complete both phases of the *join* process prior to FIFO channel reconfiguration. The first new member may never have processed the *joinagree* and *joincomit* for the second joining member prior to the second member being incorporated into the group view. When the first joining member determines its *acwnbr* and receives the external token pool, the *joincomit* token for the second member may be present. Processing the token will result in attempts to delete the corresponding nonexistent *joinagree* token never received by the first member and subsequent removal of the second member from the status table. Both of these are error conditions.

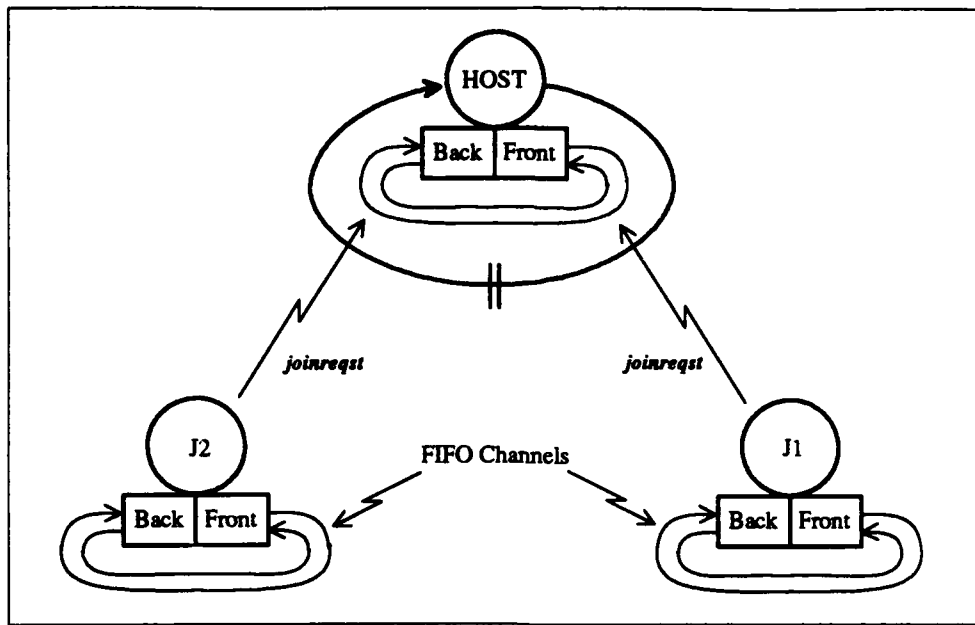


Figure 16 Simultaneous Joins

2. Solution

Initiate the FIFO reconfiguration upon transmission of the initial parameters to the new member. Figure 10, line 7.

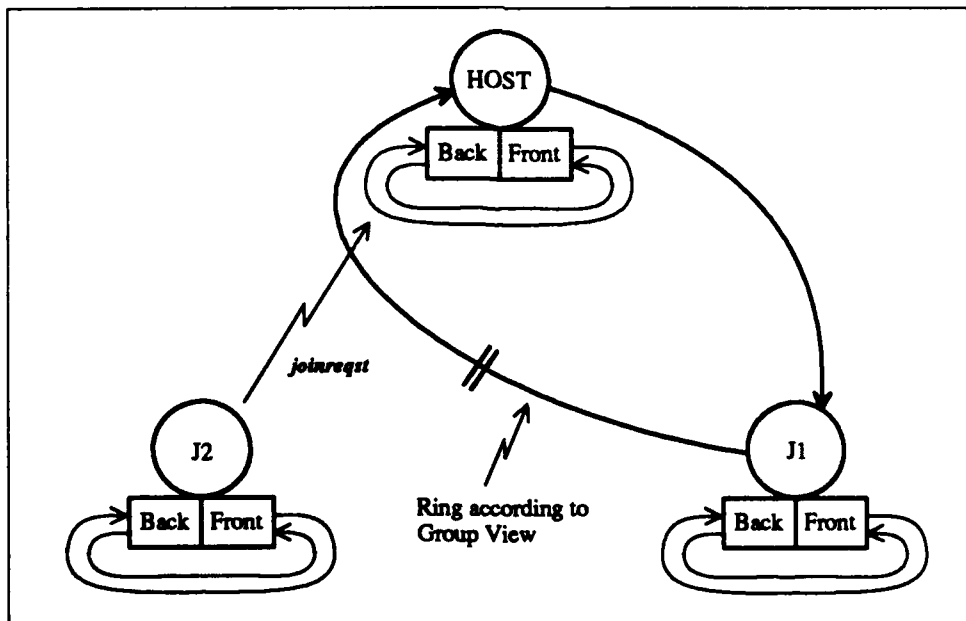


Figure 17 Group View after 1 member has joined

3. Justification

The FIFO channel can be considered to be re-configured when the host determines its *acwnbr*. The response from the new *acwnbr* is not required, as the FIFO channel reject all tokens unless they are sent by the *acwnbr*. Tokens from the old *acwnbr* are rejected and must be processed by the new *acwnbr* prior to being forwarded to the host member.




Figure 18 Group View at Host & J2

G. OTHER IMPROVEMENTS

In this section, oversights to the protocol specification and the modifications required are briefly described.

1. *Joinreqst* Token Processing

a. Problem

When processing a join message from a prospective member, *InitiateJoin* adds a *joinreqst* token to the token pool prior to generating it.

b. Solution

If the current message being processed is a join message, generate the *joinreqst* token and then add the token to the token pool. Figure 19, lines 10-13.

```
InitiateJoin for a join request message/token for  $p_{new}$  at  $p_i$ 

1  while (true)
2    if ( $p_{new} \notin ST_{p_i}, GV_{p_i}$ )
3      receive join request message or token for  $p_{new}$ 
4    end
5    if ( $p_i = p_{host}$ )
6      send initiate agreement message to AgreeProcessor for  $p_{new}$ 
7      block until AgreeProcessor acknowledges end of processing
8    else
9       $ST_{p_i}(p_{new}) \leftarrow JoinRequested$ 
10     if (join request message) /*  $p_{new}$  locates  $p_i$  and sends its join request */
11       generate  $joinreq_{p_i}(p_{new})$  token
12     end
13     add  $joinreq_{p_i}(p_{new})$  to  $TokenPool(p_i)$ 
14     send  $joinreq$  token to  $cwnbr(p_i)$ 
15   end
16 end
```

Figure 19 Processing of a Join Request Message / Token

2. Commit Token Generation

a. Problem

If a member was in a *pending* status, a commit token for that process was never generated prior to committing the change. These members would remain *pending* indefinitely.

b. Solution

Create the commit token before committing the change for a member with a *pending* status. Figure 20, lines 11-12.

ProcessCommitTkn for $commit_{p_i}(p_k)$ at p_i

```

1  if (initiate commit message received)
2    generate commit token
3    token to be processed  $\leftarrow$  generated token
4  else if (( $p_i \neq p_j$ ) && (not duplicate))
5    token to be processed  $\leftarrow$  received token
6  else
7    exit
8  end
9  CommitChange
10 while (  $p_i \in ST_{p_i}$  with pending status &  $Rank(p_i) < Rank(p_m)$ ,  $p_m \in ST_{p_i}$  )
11   generate commit token
12   token to be processed  $\leftarrow$  generated token
13   CommitChange in rank order
14 end

```

Figure 20 Generate / Receive and Process a Commit Token

3. Message Queue in the FIFO Channel Layer

The front processor was modified to transmit the head of the message queue after receiving any message, either on the internal or external channel. Figure 21.

FIFO Channel - *FRONT* Process

```
1  Wait for a channel ready to read
2  if (external channel ready)
3      if (Status_Query)
4          send Status_Query to MonitorProcess
5      else if (JoinRequest)
6          send JoinRequest to JoinProcessor
7      else if (InitialParameters)
8          send InitialParameters to JoinProcessor
9      else if (TokenAck)
10         if (Received_Serial_Number = Expected_serial_number)
11             remove Head_of_Queue
12             decrement Queue_Counter
13         end
14     end
15 else /* internal channel ready */
16     if (Token)
17         change Token to external format /* add external header */
18         insert Token in queue
19         increment Serial_Number
20         increment Queue_Counter
21     else if (TokenPool)
22         discard all messages in queue
23         change TokenPool to external format /* add external header */
24         insert TokenPool in queue
25         increment Serial_Number
26         increment Queue_Counter
27     else if (StatusReport)
28         update cwnbr
29         send StatusReport to cwnbr
30     end
31 end
32 if (Queue_Counter > 0)
33     send Head_of_Queue to cwnbr
34     set Expected_serial_number = Head_of_Queue_serial_number
35 end
```

Figure 21 FIFO Channel - Front Process

H. SYNOPSIS

This chapter has described the changes that were required to successfully implement the group membership protocol. Changes covered coding as well as protocol related problems not discovered in the original specification. These changes deal with the correct functioning of the protocol and do not address performance issues. Performance is dealt with in Chapter IV.

IV. PERFORMANCE OF THE GMP

The performance of the protocol on the Electrical and Computer Engineering Local Area Network (ECE LAN) consisting of SUN2 workstations connected via an Ethernet, was measured and the results are presented in this chapter.

A. LATENCY

The latency involved in processing changes to the group view is measured by each member using the local time clock on each specific processor. Timestamps were generated for the conditions listed in Table 2.

Table 2 CONDITIONS WARRANTING A TIME STAMP

<u>Time Stamp</u>	<u>Where Taken in Code</u>
t_{ai}	initiate an <i>agree</i> token
t_{ar}	receive an <i>agree</i> token
t_{cs}	send a <i>commit</i> token

The timestamps and related data were dumped to a file local to each processor. A filter program was designed and written to compile the data from the various systems, given a list of the members in the test group, and the maximum number of members. The filter program was designed for a restricted set of all possible group views given the above data. The case when the group view starts as the initial member, grows to the maximum number of members and then shrinks to the original host is the only possible case handled by the filter. Figure 22 illustrates the required group view changes.




Figure 22 Group Changes Required by the Filter Program

A further restriction is that all members departing the group must be the host's *acwnbr*. This removes the need for a local database in the filter program to track which members are part of the current iteration of timestamp evaluations. Future improvements to the filter program can implement a dynamic database to account for all possible changes to the test group. The format of the output file is shown in Figure 23.

agree	sp	sec	sp	usec	init						
agree	sp	sec	sp	usec	recv	change to group size	sp	subject	sp	originator	
comit	sp	sec	sp	usec	send						

sp = space

Figure 23 Time Stamp File Format

The time required to implement a change at a given member i is calculated by subtracting the initial time stamp from the completion time stamp.

$$t_i = t_{cs} - t_{ar} \mid i \neq host$$

$$t_i = t_{cs} - t_{ai} \mid i = host$$

The average time to commit a change at each member was calculated as follows:

$$\bar{t} = \frac{\sum_{i=1}^n t_i}{n}$$

However, as the group increases from n members to $n+1$ members, the *joinagree* must be processed by the n members currently in the group. Similarly, the n members

must receive and process the *joincomit* token. Hence, for a join to a group of n members, the processing time must be $n(t_{agree} + t_{comit})$, where t_{agree} is the time required to process an *agree* token and t_{comit} is the time to process a *commit* token. The communication time to transmit the token from one member to its neighbor is t_{comm} . Since each of the tokens must be transmitted to the n members, the total communication cost is $t = 2n * (t_{comm})$. Therefore, the total time required to implement a *join* at all members of a n member group is

$$t = n(t_{agree} + t_{comit} + 2t_{comm})$$

Notice that the time to implement a change is proportional to the size of the group.

Table 4 PROCESSING TIME VALUES

<u>Time</u>	<u>Occurrence</u>
t_{agree}	time to process an agree token
t_{comit}	time to process a commit token
t_{comm}	inter-member communication time

Similarly, the expected time for a failure can be determined to have a linear relationship to the size of the remaining group.

B. TESTING

The performance of the GMP was tested on the ECE LAN consisting of SUN workstations linked via an Ethernet. There were no gateways between any of the members of the group. Only single complete changes were allowed at any given time. A complete reconfiguration included the underlying FIFO channel as well as the logical ring structure. A linear relationship was observed between the number of members in the group and the average time it took to commit the new member. Figure 24. Since the communication depends heavily upon the network load as well as the individual processor load, average values over a large variety of conditions such as time of day and number of people on the network were generated in order to get reliable data points.




Figure 24 Average Time for each Member to Implement a Join to the Group

Similarly, the time required to remove a member from the group view was obtained and plotted. Thus, the relationship between time and group size was determined for a decreasing group size. Again, only single complete changes were allowed. Figure 25. A linear relationship was observed as expected.




Figure 25 Average Time to Implement a Failure in the Group

V. METHODS FOR IMPROVING PERFORMANCE

In this chapter we explore two methods that might be used to further improve the overall performance of the protocol. Each could be incorporated along with the other or individually.

A. MESSAGE REDUCTION

The protocol as currently implemented has a high overhead due to the number of inter-member messages required to effect changes in the logical ring structure. In order to reduce the number of messages required for the maintenance of the logical ring, structure three methods are discussed below.

1. TokenPool versus Tokens

Consider multiple near simultaneous changes to the group view. Transmitting the token pool instead of individual tokens will result in a reduction of messages if the changes occur close enough together such that the token pool for *one change* includes the tokens for the subsequent changes. This will result in a decreased number of messages. However, the probability of such changes occurring is minimal. Since changes to the group are uncorrelated, the probability of such changes occurring is minimal. The corresponding reduction of message traffic is negligible. Additionally, there is an increase in the size of the message for most traffic. Modifications required to effect this include the dynamic generation of the group view by the FIFO channel. Instead of receiving the token pool for generation, the FIFO channel would receive a flag and, at that point, generate the external token pool message. One such method might be to set a flag that indicates that the token pool must be transmitted. FIFO properties are maintained by the order in which the tokens are processed upon receipt. Reduction will occur only if multiple changes occur prior to the transmission of the token pool for the initial change. Problems arise in the correct setting of the transmit flag; i.e., did change #2 get sent in the last token pool,

or is another token pool transmit required. This method is not recommended since the number of messages is reduced only in special cases.

2. Periodic Token Pool

Recall that tokens are generated only if a change to the group view occurs. Another method that will reduce the number of inter-member messages is to periodically send the token pool instead of individual tokens. In this manner, multiple tokens can be transmitted simultaneously. Message reduction is indicated only if multiple changes to the group view occur within the period of the token pool transmission. If this does not occur, message traffic will actually increase; i.e., if there are no changes within this period the token pool is still transmitted. This method will also increase the latency in phase completion as tokens are not immediately forwarded around the ring. Additionally, the message size will be increased. This method is not recommended either.

3. Piggyback the Token Pool

A further refinement would be to include the local token pool as part of the status report. The monitoring member, upon receipt of a *statusrpt*, would parse the token pool and process the appropriate tokens. FIFO channel requirements are maintained by the order in which tokens are processed upon receipt of a token pool. This, however, leads to additional processing for every status report.

Difficulties might occur in the latency of cycle completion in a large group. Most notably, consider when a new member has requested to join an existing group. Define t_l as the latency within a process. It is the difference in time between receiving the token pool via a *statusrpt* and transmitting the tokens around the ring. t_l can be modeled as a random variable. Ignoring the communication time, t_{comm} , and the processing time, t_{proc} , associated with the normal processing of tokens, the latency for a change to an N member group (i.e. a join request) becomes:

$$T_{total} = 2(N \times t_l)$$

Thus, the latency involved may become prohibitive for large N. This method is recommended for implementation, provided that the latency of changes is not important.

Care must be taken to ensure the time-out on a join request is large enough to encompass the worst case scenario. However, since the time-out is finite, this method place an implicit upper bound on the maximum group size. Once the group is large enough, new members will be unable to join due to the time needed to implement the *join* around the ring.

B. SINGLE-THREADED PROGRAM

1. Problem

The current design of the protocol involves concurrent processes handling specific areas of responsibility. Fully implementing this design would allow each process to reside on different processors. However, in most cases, a single processor is the norm. There is a significant amount of overhead due to the context switching and intra-member messages. The asynchronous nature of the protocol has several areas requiring process blocking as mentioned in chapter III.

2. Solution

Redesign the main process using a single-threaded program . Figure 26 shows the recommended processes and inter-dependencies.




Figure 26 Single-Threaded Process Inter-Dependencies

3. Justification

The *TokenProcessor* would be a single-threaded program combining all aspects of the current design as shown in Table 4. The watchdog timer must still be a separate

entity by definition. The FIFO channel is not included in the *TokenProcessor* as *Back* and *Front*, are by nature, separate programs without any timing restrictions.

Table 4 SINGLE THREADED PROCESSES AND EQUIVALENTS

<u>Single-Threaded GMP</u>	<u>Multi-Threaded GMP</u>
<i>Timer</i>	<i>Timer</i>
	<i>JoinProcessor</i>
	<i>AgreeProcessor</i>
	<i>ComitProcessor</i>
	<i>IntegrateMember</i>
<i>TokenProcessor</i>	<i>Status Table Manager</i>
	<i>Group View Manager</i>
	<i>Token Pool Manager</i>
	<i>StatusMonitor</i>
	<i>StatusReporter</i>
<i>BACK</i>	<i>BACK</i>
<i>FRONT</i>	<i>FRONT</i>

The asynchronous nature of the design would be eliminated. The need for block and wait would be eliminated if a single-threaded program design were to be used. Additionally, the design would result in a significant decrease in the overhead costs due to the inter-process messages and connecting services being eliminated. The single-threaded program also eliminates the need for separate database managers. The *TokenProcessor* can maintain all databases internally, with different pointers keeping the different databases separate.

Concurrent with the new design, a review of all subroutines is recommended. Current design and programming practices preserves all data passed to the subroutines. This can lead to significant overhead since the data is stored multiple times.

VI. CONCLUSIONS AND RECOMMENDATIONS

In this thesis, the modifications required to implement the group membership protocol as proposed by [5] are presented. The protocol has been successfully implemented and to date has run continuously for more than 48 hours with a stable group membership. Additionally, a group size of 20 members was achieved. These results, though preliminary, are the first for this protocol. Although the protocol is functioning, continued debugging and improvement are currently going on.

As expected, the time required to implement a change was found to have a linear relationship to the eventual group size.

Further work should include the re-design of the protocol as a single-threaded program. In this manner, the response of the protocol can be improved as inter-process communication time is reduced drastically. However, attempting to implement the message reduction schemes to improve performance is not recommended as there is little to gain in the number of messages. On the contrary, implementation of the message reduction schemes would result in a large increase in the latency of changes to the membership.

Additional research is suggested in the area of network partitioning. Consider that a network may partition in two separate halves that are fully connected on either side of the boundary. A group originally existing on both sides will become two groups with the same name operating independently on either side of the partition. The difficulty arises when the network is repaired. The protocol does not provide for the possibility of merging the two groups back into the original group. The problem of handling network partitioning is non-trivial.

LIST OF REFERENCES

- [1] A. Ricciardi and K. Birman, "Using process groups to implement failure detection in asynchronous environments," in *ACM Symposium on Principles of Distributed Computing, Montreal, Quebec, Canada*, pages 341-353, August 1991. Also available as TR91-1188, Dept. of Computer Science, Cornell University.
- [2] Kenneth P. Birman, "The process group approach to reliable distributed computing," Technical Report TR91-1216, Cornell University Computer Science Department, Ithaca, NY, July 1991.
- [3] S. B. Shukla, F. Pires, and D. Raghuram, "Design Implementation and Performance of a Decentralized Group Membership Protocol for Asynchronous Environments Using Ordered Views," Technical Report NPS-EC-93-006, Naval Postgraduate School, Monterey, California
- [4] Shridhar B. Shukla and Devalla Raghuram, "Group Membership in Asynchronous Distributed Environments Using Logically Ordered Views," Technical Report NPS-EC-92-009, Naval Postgraduate School, Monterey, California
- [5] Fernando Pires, "Design of a Decentralized Asynchronous Group Membership Protocol and an Implementation of Its Communications Layer," Master's Thesis, March 1993, Naval Postgraduate School, Monterey, California
- [6] Flaviu Cristian, "Agreeing on who is present and who is absent in a synchronous distributed system," in *Proceedings of the 18th International Conference on Fault Tolerant Computing, Tokyo, Japan*, pages 206-211, 1988.
- [7] W. Richard Stevens, *Unix Network Programming*, Prentice Hall, 1990.

APPENDIX

The following code is included for completeness.

TABLE OF CONTENTS

Definitions & Utilities	42
Simple Application & Main Process	85
FIFO Channel Processes	95
Monitor Processes	105
Agree Processor	116
Commit Processor	131
Integrate Member Process	144
Join Processor	151
Database Managers	159
Data Cruncher Program	180

DEFINITIONS

and

UTILITIES

Definitions	43
FIFO Utilities	44
GMP Utilities	50
Message Utilities	67
Socket Utilities	77

```

/*
 * Definitions for GMP programs
 */

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sys/ocul.h>
#include <signal.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <sys/un.h>
#include <netdb.h>
#include <sys/uio.h>
#include <errno.h>

#define CALLOC(n,type) (type *)calloc((unsigned) n,sizeof(type))
#define REALLOC(ptr,n,type) (type *)realloc(ptr,(unsigned) (n*sizeof(type)))

#define FRONT_PORT 5432
#define BACK_PORT 5433
#define SERV_HOST_ADDR "131.120.20.102" /* host address for server sun2 */
#define UNIXSTR_PATH "/s.unixstr"
#define UNIXSTR_TMPL "/tmp/so.XXXXXXX"
#define MAXLINE 255
#define TRUE 1
#define FALSE 0
#define MAXPATH 15
#define MAXPORT 6
#define MAXFD 6
#define MAXNUM 6

/* The following definitions do not account for the terminating NULL */
#define MAXHOSTNAME 50
#define MAXFNAME 15
#define MAXLIMITSIZE MAXIPNAME + 2*MAXPORT + 2
#define BUFSIZE 512
#define NODATAMSG 10

```

```

#define TIMERMSG 16
#define HEADERSIZE 11

#ifdef INADDR_NONE
#define INADDR_NONE 0xffffffff
#endif /* INADDR_NONE */

#define INVALIDMSG 0
#define TOKENOKN 1
#define TOKENPOOL 2
#define TOKENACKN 3
#define STATUSQRY 4
#define STATUSRPT 5
#define STATUSTBL 6
#define INITPARAM 7
#define JOINREQST 8
#define UPDSTATUS 9
#define GROUPVIEW 10
#define UPDATVIEW 11
#define SNDINIPAR 12
#define INITOKEN 13
#define VIEWREQST 14
#define STATREQST 15
#define TOKPREQST 16
#define INITGVIEW 17
#define INITTABLE 18
#define INITPOOL 19
#define TIMEOUT__ 20
#define STARTTMR 21
#define DELTOKEN 22
#define EXTKNPOOL 23

```

Table of Contents

FIFO Utilities	45
enqueue	45
dequeue	46
get_queue_head	46
flush_queue	47
send_msg_front	47
send_msg_back	48
send_ack	48
send_msg_in	49

```

/*****
* QUEUE HANDLING AUXILIARY FUNCTIONS FOR FIFO PROCESSES
* (FIFOUTIL.C)
*
* The following functions are available to be used:
* void enqueue(queue *quptr, char *msg);
* void dequeue(queue *quptr);
* void get_queue_head(queue *quptr, char **msg);
* void flush_queue(queue *quptr);
* void send_ack(char *msg, char *orig, char *dest);
* void send_msg_back(char *msg, char *dest);
* void send_msg_front(char *msg, char *dest);
* void send_msg_in(char *msg, char *dest);
*
* Refer to the function header comments for detailed info.
* Some functions in this file need socutil.c and msgutil.c
*
*****/
* Written by:  Fernando J. Pires
*             David Pezdirtz
*
* Last revision:  12 Jul 1993
*
*****/

```

```

/*****
enqueue - inserts the external message 'msg' at the
tail of the queue 'quptr'.
The original 'msg' is not modified.
'quptr' must be created before the first call to
enqueue(). To create a queue use:

queue qu;           / declaration /
queue *quptr = &qu; / initialize pointer /
qu.tail = qu.head = NULL; / empty queue /

enqueue(quptr, msg); / function call /
*****/
void enqueue(head, tail, msg)
link **head, **tail;
char *msg;
{
    link *pntmnt, *tmp;
    pntmnt = CALLOC(1, link);
    pntmnt->data = CALLOC(strlen(msg) + 1, char);
    strcpy(pntmnt->data, msg);
    pntmnt->next = NULL;

    if (*head == NULL) {
        *head = pntmnt;
        *tail = pntmnt;
    }
    else {
        tmp = *tail;
        tmp->next = pntmnt;
        *tail = pntmnt;
    }
} /* end enqueue */

```



```

*****
dequeue - remove a msg from the head of the queue 'qptr'.

Sample call: dequeue(qptr);

*****
void dequeue(head, tail)
link **head, **tail;
{
    link *tmp;
    if (*head != NULL) {
        tmp = *head;
        *head = tmp->next;
    }
    if (head == NULL) {
        *tail = NULL;
    }
    free(tmp->data);
    free(tmp);
} /* end if qptr */

} /* end dequeue */

```

```

*****
get_queue_head - returns a pointer to the message at the head of the queue.

Sample call: get_queue_head(qptr, &msg);

*****
void get_queue_head(head, msg)
link *head;
char **msg;
{
    if (head == NULL) {
        *msg = NULL;
    }
    else {
        *msg = head->data;
    } /* end get_queue_head */
}

```

```

/*****
flush_queue - remove all nodes of 'qpur' from memory.
All used memory is deallocated.
'qpur' remains a valid empty queue and can be
reused by enqueue().
*****/
void flush_queue(head, tail)
link **head, **tail;
{
    link *tmp;

    while (*head != NULL) {
        tmp = *head;
        *head = tmp->next;

        free(tmp->data);
        free(tmp);
    }

    *tail = NULL;
}

/* end flush_queue */

```

```

/*****
send_msg_front - sends an external message 'msg' to the front port of the
specified IP destination 'dest'. 'dest' is a string with element address format.
The original message is not disturbed.

Sample call: send_msg_front(msg, dest);
*****/
void send_msg_front(msg, dest)
char *msg, *dest;
{
    link *list;
    char *IPaddr, *frontport, *tmp;
    u_short port;

    /* make a copy of the address */
    tmp = CALLOC(strlen(dest) + 1, char);
    strcpy(tmp, dest);

    list = str2list(tmp, ";");

    if (getfromlist(list, &IPaddr, 1) != 1) {
        printf("send_msg_front: IP address error\n");
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1);
    }

    if (getfromlist(list, &frontport, 2) != 2) {
        printf("send_msg_front: front port error\n");
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1);
    }

    /* convert port # to network format */
    port = htons((u_short)atoi(frontport));

    sendmsg(msg, strlen(msg), IPaddr, port);

    remove1st(list);
    free(tmp);
} /* end send_msg_front */

```


Table of Contents

GMP Utilities	51
GetAcwnbr	52
InStatusTable	55
CountDown	56
SendTkn2Agr	56
InGroup	57
GetMembWithRank	58
InTokenPool	59
TokensREqual	60
GetRank	61
GetTokenType	62
GetGroupSize	62
GetStatus	63
RelativeRank	63
GetGroupView	64
GetStatusTable	64
GetTokenPool	65
agreetoken	65
committoken	65
first_time	66

* Utilities for the protocol processes. (gmputil.c)

* Description: GetAcwnbr0;

* InStatusTable0;

* InGroup0;

* CountDown0;

* SendTtn2Agr0;

* InTokenPool0;

* TokensREqual0;

* GetRank0;

* GetTokenTyped0;

* GetGroupSize0;

* GetMembrWithRank0;

* GetStatus0;

* RelativeRank0;

* GetGroupView0;

* GetStatusTable0;

* GetTokenPool0;

* agreeToken0;

* commitToken0;

* first_time0;

* Written by: Shridhar Shukla

* David Pezdirtz

* Date: 23 Nov 1993

#define FAILAGREE 0

#define FAILCOMMIT 1

#define JOINAGREE 2

#define JOINCOMMIT 3

#define JOINRQSTT 4

#define NONBLOCKING 1

#define BLOCKING 0

#define GVLSTOFFSET 4 /* 4th field is the first element in gv list */

#define STLSTOFFSET 3 /* 3rd field is the first element in st list */

#define TKPLISTOFFSET 3 /* 3rd field is the first element in token pool list */

#define COUNT_DOWN_STEP 1000 /* 1 timer tick is AT LEAST 1000 micro seconds. */

#define RESENDREQST 10000 /* resend request after at least 10 seconds. */

#define TPAD_INTERVAL 250 /* 25 sec */

#define TQRY_INTERVAL 500 /* 5 sec */

/* The following definitions do not account for the terminating NULL. */

#define MSGTYPELEN 9

#define TOKENTYPELEN 9

#define STSTYPELEN 9

#define QUERYLEN MSGTYPELEN+2*MAXMLMTSIZE+3

#define INITTOKENMSGLEN 2*MSGTYPELEN+2+MAXMLMTSIZE+1

#define INITGVMSGLEN MSGTYPELEN+2*MAXXNUM+MAXMLMTSIZE+4

#define TOKENLEN 2*MAXMLMTSIZE+TOKENTYPELEN+3

#define TOKENMSGLEN TOKENLEN+MSGTYPELEN+2

#define UPDSTMSGLEN MSGTYPELEN+MAXMLMTSIZE+STSTYPELEN+3

#define UPDVIEWLEN 14+MAXMLMTSIZE+1

#define DELTKNLEN 10+TOKENLEN+1

#define SNDINPLEN 10+MAXMLMTSIZE+1


```

close(gvnmfd);

/*acquire status table */
strcpy(request, "sareqs#");
msglen = strlen(request);
smfd = connectUN(smsoc);

if ( ( writensg(smfd, request, msglen) < msglen) ) {
    printf("\nGetAcwnbr: reqst to sm failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

if ( ( msglen==readmsg(smfd, &stbuf, "#") < 0 ) ) {
    printf("\nGetAcwnbr: st read failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

stbuf[msglen] = NULL;

close(smfd);

/*
 * make less from the message buffers and get list sizes
 * for group view  and for status table
 */
st = str2list(stbuf, "\n=#");

if ( getfromlist(st, &ssize, 2) == 0 ) {
    printf("\nGetAcwnbr: st msg parsing for size failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

ssize = atoi(ssizestrng);
removelis(st);

/*
 * Search for myaddr in gv and return the element preceeding it
 * (modulo group size) that is NOT part of the status table.
 *
 */
myrank = GetRank(gvbuf, myaddr);

if ( myrank < 0 ) {
    printf("\nGetAcwnbr: parse on myaddr search failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

gv = str2list(gvbuf, "\n=#");

if ( getfromlist(gv, &gvsizestrng, 3) == 0 ) {
    printf("\nGetAcwnbr: gv msg parsing for size failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

gvsizestrng = atoi(gvsizestrng);

if (myrank == 0)
    candidateposition = gvsizestrng + GVLSTOFFSET - 1;
else
    candidateposition = myrank + GVLSTOFFSET - 1;

if ( getfromlist(gv, &candidate, candidateposition) == 0 ) {
    printf("\nGetAcwnbr: parse on nbr search failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

/*
 * Check st for every group member anticlockwise from myaddr
 */
found = FALSE; /* found the Acwnbr */
searchcomplete = FALSE; /* finished looking but did not find one */

```



```

/* all others in st */
if ( candidateposition == (myrank + GVLISTOFFSET) ) {
    searchcomplete = TRUE;
    strcpy(acwnbr, myaddr);
}
else
    if ( getfromlist(gv, &candidate, candidateposition) == 0 ) {
        printf("\nGetAcwnbr: parse on nbr search failed.\n");
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    } /* end else InStatusTable */
} /* end while found */

removelst(gv);
free(gvbuf);
free(stbuf);

if (found != searchcomplete)
    return(0);
else {
    printf("\nGetAcwnbr: expected flow of execution failed.\n");
    printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1);
} /* end GetAcwnbr */

```

```

if (ssize == 0) {
    found = TRUE;
    strcpy(acwnbr, candidate);
}

while ( (found == FALSE) && (searchcomplete == FALSE) ) {
    free(stbuf);

    /* acquire status table */
    strcpy(request, "statreqst#");
    smsglen = strlen(request);
    smid = connectUN(smsoc);

    if ( ! writmsg(smid, request, smsglen) < smsglen ) {
        printf("\nGetAcwnbr: reqst to sm failed.\n");
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    if ( (smsglen=readmsg(smid, &stbuf, "#") < 0 ) ) {
        printf("\nGetAcwnbr: st read failed.\n");
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    stbuf[smsglen] = NULL;

    close(smid);

    if ( ! InStatusTable(stbuf, candidate, (char *) NULL) == FALSE ) {
        /* candidate is acwnbr if it is not in st */
        found = TRUE;
        strcpy(acwnbr, candidate);
    }
    else { /* rotate ring position */

        if ( candidateposition == GVLISTOFFSET )
            candidateposition = gsize + GVLISTOFFSET - 1;

        else
            candidateposition--;
    }
}

```

```

/*****

```

```

InStatusTable: returns TRUE if candidate is in the status table, FALSE if not
Initializes status to the value from the status table unless the address
passed as status is NULL,

```

```

When status is not required: InStatusTable(stbuf, candidate, (char *) NULL)

```

```

When status is required: InStatusTable(stbuf, candidate, status)

```

```

where status is the address of a memory
allocated null terminated string of 10
characters including the null character.

```

```

*****/

```

```

int InStatusTable(stbuf, candidate, status)

```

```

char *stbuf, *candidate, *status;

```

```

{

```

```

    int nextinst, ssize, found;

```

```

    char *smember, *ssizestring, *ss, *stble;

```

```

    link *st;

```

```

    /* copy the status table */

```

```

    stble = CALLOC(strlen(stbuf) + 1, char);

```

```

    strcpy(stble, stbuf);

```

```

    /* separate elements from their status */

```

```

    st = str2list(stble, "\n= #");

```

```

    if ( getfromlist(st, &ssizestring, 2) == 0 ) {

```

```

        printf("\nStatusTable: st msg parsing for size failed\n");

```

```

        printf("\n07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");

```

```

        exit(-1);
    }

```

```

    ssize = atoi(ssizestring);

```

```

    if (ssize == 0)

```

```

        found = FALSE;

```

```

    else { /*search the status table entries */

```

```

        found = TRUE;

```

```

        nextinst = STLSTOPSET; /*first member in st is the third field */

```

```

    if ( getfromlist(st, &smember, nextinst) == 0 ) {
        printf("\nStatusTable: first parse on st failed\n");
        printf("\n07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1);
    }

```

```

    while ( (found == TRUE) && (strcmp(smember, candidate) != 0) ) {
        nextinst += 2; /* skip the status field */

```

```

        if ( getfromlist(st, &smember, nextinst) == 0)
            found = FALSE;

```

```

    } /* end while found */

```

```

    if (status != NULL) {

```

```

        if (found == TRUE) {

```

```

            if (getfromlist(st, &sts, nextinst + 1) == 0) {
                printf("\nStatusTable: initializing status failed\n");
                printf("\n07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
                exit(-1);
            }

```

```

            strcpy(status, sts);

```

```

        } /* end if found */

```

```

    } /* end if status */

```

```

    } /* end else ssize */

```

```

    removelist(st);

```

```

    free(stble);

```

```

    return(found);

```

```

} /* end InStatusTable */

```

```

/*****
CountDown: Decrements the location at valptr after at least
'step' microseconds.
*****/
int CountDown(valptr, step)
int *valptr, step;
{
    usleep(step);
    return(*valptr = *valptr - 1);
} /* end CountDown */

```

```

/*****
SendTkn2Agr: A token is sent to the agreement process and the caller blocks
until the agreement process acks.
        exits upon failure to send.
*****/
void SendTkn2Agr(agsoc, tkntype, tknsubject)
char *agsoc, *tkntype, *tknsubject;
{
    int agrfd, msglen;
    char c[1], initagree[INITTOKENMSGLEN + 1];

    /* assemble initiate agreement msg */
    strcpy(initagree, "initoken\n"); /* header */
    strcat(initagree, tkntype); /* type */
    strcat(initagree, ""); /* separator */
    strcat(initagree, tknsubject); /* subject to token */
    strcat(initagree, "#"); /* end of msg */

    /* send initiate agreement msg and block */
    msglen = strlen(initagree);
    agrfd = connectUN(agsoc);

    if ( writemsg(agrfd, initagree, msglen) < msglen ) {
        printf("\nSendTkn2Agr: init agree send failed for %s\n", tkntype);
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1);
    }

    if ( read(agrfd, c, 1) < 0 ) {
        printf("\nSendTkn2Agr: empty msg read failed after sending %s\n", tkntype);
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1);
    }

    close(agrfd);
} /* end SendTkn2Agr */

```

```

/*****
InGroup: returns the rank of the candidate from the group view. -1 if
the candidate is not found.
*****/
int InGroup(gvmsoc, candidate)
char *gvmsoc, *candidate;
{
    int found, gvmfid, msglen, rank, gvsiz;
    char *gvbuf, *element, request[NODATAMSG + 1], *gvsizstring;
    link *gv;

    /*acquire group view */
    strcpy(request, "viewreqst#");
    msglen = strlen(request);
    gvmfid = connectUN(gvmsoc);

    if ( writemsg(gvmfid, request, msglen) < msglen) {
        printf("\nInGroup: reqst to gvm failed\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    if ((msglen=readmsg(gvmfid, &gvbuf, "#") < 0) {
        printf("\nInGroup: gv read failed\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    gvbuf[msglen] = NULL;
    close(gvmfid);

    /*
    * make a list from the message buffer and get the host
    */
    gv = str2list(gvbuf, "\n#");

    if ( getfromlist(gv, &gvsizstring, 3) == 0 ) {
        printf("\nInGroup: msg parsing for gv size failed\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }
}

gvsiz = atoi(gvsizstring);
rank = 0;
found = FALSE;

while ( (found == FALSE) && (rank <= gvsiz-1) ) {
    if ( getfromlist(gv, &element, rank + 4) == 0) {
        printf("\nInGroup: gv parsing for element rank %d failed\n", rank);
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    if ( strcmp(element, candidate) == 0)
        found = TRUE;
    else rank++;
} /* end while found */

removelis(gv);
free(gvbuf);

if ( found == TRUE )
    return(rank);
else return(-1);

} /* end InGroup */

```



```

*****
InTokenPool: returns TRUE if token is in the token pool, FALSE if not.
For token equality, only checks the subject and the type,
and ignores the originator.
*****
int InTokenPool(tokenpool, token)
char *tokenpool, *token;

{
    int found, next, done;
    char *candidate, *poolsizestring, *tp;
    link *tqlist;

    /*
     * make a list from the message buffer and get the pool size
     */
    tp = CALLOC(strlen(tokenpool) + 1, char);
    strcpy(tp, tokenpool);

    tqlist = str2list(tp, "\n-#");

    found = FALSE;
    done = FALSE;

    next = TKPLUSOFFSET; /* first member in token pool is the third field */

    while ( done == FALSE){
        if ( getfromlist(tqlist, &candidate, next) == 0)
            done = TRUE;
        else{
            if (TokensREqual(token, candidate) == TRUE) {
                done = TRUE;
                found = TRUE;
            }
            next++;
        }
    }

    } /* end while !done */

```

```

removelst(tqlist);
free(tp);

return(found);
} /* end InTokenPool */

```



```

/*****
GetTokenType: does not disturb the token passed in.
returns -1 if the token type is invalid.
*****/
int GetTokenType(token)
char *token;
{
    int type;
    char token[TOKENLEN + 1], *tktype;
    link *tknl;

    strcpy(tknl, token);
    tknl = str2list(tknl, " ");

    if (getfromlist(tknl, &tktype, 1) == 0) {
        printf("\nGetTokenType: local token parsing for type failed.\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1);
    }

    type = -1;

    if (strcmp(tktype, "failagree") == 0)
        type = FAILAGREE;
    if (strcmp(tktype, "failcommit") == 0)
        type = FAILCOMIT;
    if (strcmp(tktype, "joinagree") == 0)
        type = JOINAGREE;
    if (strcmp(tktype, "joincommit") == 0)
        type = JOINCOMIT;
    if (strcmp(tktype, "joinreqst") == 0)
        type = JOINRQSTT;

    removelist(tknl);

    return(type);

} /* end GetTokenType */

```

```

/*****
GetGroupSize: returns the group view size. Requires the caller to supply a
string containing the group view.
*****/
int GetGroupSize(gv)
char *gv;
{
    int gvsz, gvlen;
    char *lgv, *gvszstrng;
    link *gvl;

    gvlen = strlen(gv);
    lgv = CALLOC(gvlen+1, char);
    strcpy(lgv, gv);

    /*
    * make a list from the local copy
    */
    gvl = str2list(lgv, "\n==");

    if (getfromlist(gvl, &gvszstrng, 3) == 0) {
        printf("\nGetGroupSize: msg parsing for gvsz failed.\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1);
    }

    gvsz = atoi(gvszstrng);
    removelist(gvl);
    free(lgv);

    return(gvsz);

} /* end GetGroupSize */

```

```

/*****
GetStatus: returns the status of the candidate from the status table. -1 if
the candidate is not found. Requires the caller to supply a string
containing the status table.
*****/
int GetStatus(st, candidate)
char *st, *candidate;
{
    int stat;
    char *status;
    stat = -1;
    status = CALLOC(9 + 1, char);

    /* if subject is in status table */
    if (InStatusTable(st, candidate, status) == TRUE) {
        if ( strcmp(status, "failagree") == 0)
            stat = FAILAGREE;
        if ( strcmp(status, "failcommit") == 0)
            stat = FAILCOMMIT;
        if ( strcmp(status, "joinagree") == 0)
            stat = JOINAGREE;
        if ( strcmp(status, "joincommit") == 0)
            stat = JOINCOMMIT;
        if ( strcmp(status, "joinqstd") == 0)
            stat = JOINRQSTT;
    }

    free(status);
    return(stat);
} /* end GetStatus */

```

```

/*****
RelativeRank: returns the relative rank of the subject wrt origin. Modulo
gv_size. The relative rank of the original process is zero. Returns
(-1) if either subject, or origin are not in group view.
*****/
int RelativeRank(gv, subject, origin)
char *gv, *subject, *origin;
{
    int rank_s, rank_o, rr;
    rank_s = GetRank(gv, subject);
    rank_o = GetRank(gv, origin);

    /* calculate the rr (assuming subject and origin in gv) */
    rr = rank_s - rank_o;

    if (rr < 0) {
        rr = rr + GetGroupSize(gv);
    }

    /* check for error condition */
    if ((rank_s < 0) || (rank_o < 0))
        return(-1);
    else
        return (rr);
} /* end RelativeRank */

```

```

/*****

```

GetGroupView: returns a pointer to the group view. Requires gvmsock as a parameter.

Typical call:

```
buf = GetGroupView(gvmsock);
```

```
...
```

```
free(buf);
```

```

*****/

```

```
char *GetGroupView(gvmsock)
```

```
char *gvmsock;
```

```
{
```

```
    int msglen, gvmfd;
```

```
    char gvmreq[NODATAMSG + 1], *gv;
```

```
    /* get local group view */
```

```
    strcpy(gvmreq, "viewreq#");
```

```
    msglen = strlen(gvmreq);
```

```
    gvmfd = connectUN(gvmsock);
```

```
    if ( writemsg(gvmfd, gvmreq, msglen) < msglen ) {
```

```
        printf("GetGroupView: group view request failed.\n");
```

```
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
```

```
        exit(-1);
```

```
    if ( (msglen=readmsg(gvmfd,&gv,"#")) < 0 ) {
```

```
        printf("GetGroupView: group view read failed.\n");
```

```
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
```

```
        exit(-1);
```

```
    gv[msglen] = NULL;
```

```
    close(gvmfd);
```

```
    return(gv);
```

```
}
```

```

/*****

```

GetStatusTable: returns a pointer to the status table. Requires srmsock as a parameter.

Typical call:

```
buf = GetStatusTable(srmsock);
```

```
...
```

```
free(buf);
```

```

*****/

```

```
char *GetStatusTable(srmsock)
```

```
char *srmsock;
```

```
{
```

```
    int msglen, srmfd;
```

```
    char srmreq[NODATAMSG + 1], *st;
```

```
    /* get status table */
```

```
    strcpy(srmreq, "srmreqs#");
```

```
    msglen = strlen(srmreq);
```

```
    srmfd = connectUN(srmsock);
```

```
    if ( writemsg(srmfd, srmreq, msglen) < msglen ) {
```

```
        printf("GetStatusTable: status table request failed.\n");
```

```
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
```

```
        exit(-1);
```

```
    if ( (msglen=readmsg(srmfd,&st,"#")) < 0 ) {
```

```
        printf("GetStatusTable: status table read failed.\n");
```

```
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
```

```
        exit(-1);
```

```
    st[msglen] = NULL;
```

```
    close(srmfd);
```

```
    return(st);
```

```
}
```

```

*****
GetTokenPool: returns a pointer to the token pool. Requires tpmsoc as a
parameter.

Typical call:

    buf = GetTokenPool(tpmsoc);

    ...

    free(buf);

*****
char *GetTokenPool(tpmsoc)
char *tpmsoc;
{
    int msglen, tpnfd;
    char tpreq[NODATAMSG + 1], *tp;

    /* get local token pool */
    strcpy(tpreq, "tokpreq#");
    msglen = strlen(tpreq);
    tpnfd = connectUN(tpmsoc);

    if ( write(msg, tpnfd, tpreq, msglen) < msglen ) {
        printf("GetTokenPool: token pool request failed.\n");
        printf("*****\n");
        exit(-1);
    }

    if ( (msglen = read(msg, tpnfd, "#")) < 0 ) {
        printf("GetTokenPool: token pool read failed.\n");
        printf("*****\n");
        exit(-1);
    }

    tp[msglen] = NULL;
    close(tpnfd);

    return(tp);
}

```

```

/*****
agreeoken: return TRUE if tokentype = FAILAGREE or JOINAGREE
*****
int agreeoken(token)
char *token;
{
    int tokentype;

    tokentype = GetTokenType(token);

    if ((tokentype == FAILAGREE) || (tokentype == JOINAGREE))
        return(TRUE);
    else
        return(FALSE);
} /* end agreeoken */

/*****
commitoken: return TRUE if tokentype = FAILAGREE or JOINAGREE
*****
int commitoken(token)
char *token;
{
    int tokentype;

    tokentype = GetTokenType(token);

    if ((tokentype == FAILCOMMIT) || (tokentype == JOINCOMMIT))
        return(TRUE);
    else
        return(FALSE);
} /* end commitoken */

```

```

/*****
first_time:

return FALSE if token has been processed. A token has been processed
for the following conditions:

    if join -> has been processed if subj mbr GV
    if fail -> has been processed if subj NOT mbr GV

return TRUE otherwise.

NOTE: it is not necessary to detect all possible outcomes... the token is
checked for prior processing ONLY if it is in the external token pool and not
the local token pool. This condition can only occur if: the token has never
been seen at this member, or the token has been deleted. If the token has been
deleted, the corresponding commit must have occurred. Thus, check only the final
result as intermediate stages do not effect the local token pool.

*****/
int first_time(token, gv, st)
char *token, *gv, *st;
{
    int token_type, rntval = TRUE;
    char *tmpcn, *subject;
    link *list;

    tmpcn = CALLOC(strlen(token) + 1, char);
    strcpy(tmpcn, token);
    list = str2list(tmpcn, " #");

    if ( getfromlist(list, &subject, 2) == 0 ) {
        printf("\nfirst_time: parsing failed token subj:\n");
        printf("\ntoken = %s\n", token);
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    /* get the token type */
    token_type = GetTokenType(token);

    switch (token_type) {
        case JOINRQSTT:
        case JOINAGREE:
        case JOINCOMMIT:
            if (GetRank(gv, subject) != -1) /* subj in GV */
                rntval = FALSE;
            break;

        case FAILAGREE:
        case FAILCOMMIT:
            if (GetRank(gv, subject) == -1) /* subj NOT in GV */
                rntval = FALSE;
            break;

        default: /* error condition */
            printf("ProcessTokenPool: (first_time) invalid token type\n");
            printf("07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
            exit(1);
    }

    remove(list, list);
    free(tmpcn);
    return(rntval);
} /* end first_time */

```

Table of Contents

Message Utilities	68
str2list	68
list2str	69
removelist	70
getfromlist	70
listsize	71
int_2_ext	71
get_sr_nbr	72
msg_type	73
in_msg_type	74
ext_msg_type	74
get_target	75
get_originator	75
get_ext_target	76

```

/*****
* MESSAGE HANDLING AUXILIARY FUNCTIONS (msgutil.c)
*
* The following functions are available to be used:
*
* link *str2list(char *str, char *token);
* char *list2str(link *list, char *header, char* hock, char *tok);
* void removelist(link *list);
* int listsize(link *list);
* int getfromlist(link *list, char **str, int n);
* char *int_2_ext(char *inmsg, int startbr, char *orig);
* int get_sr_nbr(char *msg);
* int in_msg_type(char *inmsg);
* int ext_msg_type(char *extmsg);
* char *get_target(char *str);
* char *get_originator(char *str);
* char *get_ext_target(char *str);
*
* Refer to the function header comments for detailed info.
* Other functions in this file are used internally, and should
* not be used directly.
*
*****
* Written by:  Fernando J. Pires
*             David Pezdirtz
*
* Last revision:  27 Jul 1993
*
*****
struct link {
    char*data;
    struct link*next; };

typedef struct link link;

```

```

/*****
str2list - parse a string, creating a list of nodes, each
of which points to a field of the original string.
The fields are originally separated by token.
In the original string, tokens are replaced by NULL
After the list is no longer needed, removelist()
must be called for garbage collection.
*****/
link *str2list(str, token)
char *str;
char *token;
{
    link *msglst, *tmp;
    char *ptr;

    tmp = msglst = CALLOC(1, link);

    tmp->data = ptr = strtok(str, token);

    while (ptr = strtok(NULL, token)) {
        tmp->next = CALLOC(1, link);
        tmp = tmp->next;
        tmp->data = ptr; }

    return(msglst);
} /* end str2list */

```

```

*****
list2str - assembles a string from a list generated by
str2list() and appends it to the string 'header'.
'hlok' is inserted after the header. 'hlok' is
inserted in between each new field and a NULL is
added at the end.
Notes:
    'header' has to be dynamically allocated
    (it cannot be static data). The best way to
    initialize it is to use "header = CALLOC(1, char);"
    If the list is empty the header is returned
    without changes.
    The resulting message has to be deallocated
    with a free() call when it is no longer needed.
*****/
char* list2str(list, header, hlok, llok)
link *list;
char *header;
char *hlok;
char *llok;

{
    int len = 0;
    link *ptr = list;

    if (list) {
        while (ptr) { /* determine size of string to be used */
            len += strlen(ptr->data) + 1; /* Reserve space for token and NULL */
            ptr = ptr->next; }

        header = REALLOC(header, strlen(header) + 2 + len, char);

        if (hlok)
            strcat(header, hlok); /* insert hlok */

        if (len) {
            while (list) { /* assemble the string */

```

```

            strcat(header, list->data);
            list = list->next;

            if (list && llok) /*add llok except after last field*/
                strcat(header, llok);

        } /* end while list */

    } /* end if len */

    } /* end if list */

    return(header);

} /* end list2str */

```



```

/*****
    remove - Deallocates the space used by str2list() to
    generate a list. This function must be called for
    every list, once it is no longer needed.
*****/
void remove(list)
link *list;
{
    if (list->next)
        remove(list->next);
    free(list);
} /* end remove */

```

```

/*****
    getfromlist - Get the nth field from list. Upon execution 'str'
    points to the nth field.
    Returns n if the call is successful, or zero if the
    list has less than n fields.
*****/
int getfromlist(list, str, n)
link *list;
char **str;
int n;
{
    int p;

    if (list == NULL)
        return(0);

    if (n == 1) {
        *str = list->data;
        return(n);
    }
    else {
        p = getfromlist(list->next, str, n-1);

        if (p)
            return(n);
        else
            return(0);
    } /* end else n */
} /* end getfromlist */

```

```

/*****
    listsize - Return the number of elements of a list
    Returns n > 0 for a non-empty list, and 0 otherwise.
*****/

int listsize(list)
link *list;
{
    int n = 0;

    while (list) {
        list = list->next;
        n++;
    }
    return(n);
} /* end listsize */

/*****
    int_2_ext - convert 'inmsg' into an external message.
    'snr' is converted to a string and is used as
    a prefix to 'inmsg'. 'orig' is the address of the
    local element and it is inserted after 'snr'.
    An NL character is used to separate the fields.
    The original 'inmsg' is not modified.
    To convert an internal message to external format use:

        extmsg = int_2_ext(inmsg, snbr, orig);
        ...
        free(extmsg);

    If the serial number is not relevant set 'snbr' to 0.
*****/

char *int_2_ext(inmsg, snbr, orig)
char *inmsg;
int snbr;
char *orig;
{
    int msgsize;
    char temp[HEADERSIZE];
    char *extmsg;

    msgsize = strlen(inmsg);
    sprintf(temp, "%d\n" snbr);
    extmsg = CALLOC(strlen(temp) + strlen(orig) + msgsize + 2, char);
    strcpy(extmsg, temp);
    strcat(extmsg, orig);
    strcat(extmsg, "\n");
    strcat(extmsg, inmsg);
    return(extmsg);
} /* end int_2_ext */

```

```

*****
get_sr_nbr - extracts the serial number of a message that
was previously retrieved from the queue, or
received at the external port.
The original message is not disturbed.
The function returns the serial number, or -1
if an error occurs.
*****/

int  get_sr_nbr(msg)
char *msg;

(
    char *tmp, *smbstr;
    link *list;

/* make a copy of the message */
tmp = CALLOC(strlen(msg) + 1, char);
strcpy(tmp, msg);

/* break the message into a list of fields */
list = sr2list(tmp, "\n");

if (getfromlist(list, &smbstr, 1) != 1) { /* get 1st field */
    printf("get_sr_nbr error\n");
    removevlist(list);
    free(tmp);
    return(-1); }

removevlist(list);
free(tmp);
return(atol(smbstr));
} /* end get_sr_nbr */

```

```

/*****
msg_type - returns the integer value corresponding to
the type of the string 'type', as defined in grp.h
*****/

int msg_type(type)
char *type;
{
    int msgtype = INVALIDMSG;

    if (strcmp(type, "tokenackn") == 0)
        msgtype = TOKENACKN;
    if (strcmp(type, "tokenpool") == 0)
        msgtype = TOKENPOOL;
    if (strcmp(type, "tokenackn") == 0)
        msgtype = TOKENACKN;
    if (strcmp(type, "statusq") == 0)
        msgtype = STATUSQ;
    if (strcmp(type, "statusrpt") == 0)
        msgtype = STATUSRPT;
    if (strcmp(type, "statusbl") == 0)
        msgtype = STATUSBL;
    if (strcmp(type, "iniparam") == 0)
        msgtype = INITPARAM;
    if (strcmp(type, "joinreqst") == 0)
        msgtype = JOINREQST;
    if (strcmp(type, "updstatus") == 0)
        msgtype = UPDSTATUS;
    if (strcmp(type, "groupview") == 0)
        msgtype = GROUPVIEW;
    if (strcmp(type, "updateview") == 0)
        msgtype = UPDATVIEW;
    if (strcmp(type, "sndinipar") == 0)
        msgtype = SNDINIPAR;
    if (strcmp(type, "initoken") == 0)
        msgtype = INITOKEN;
    if (strcmp(type, "viewreqst") == 0)
        msgtype = VIEWREQST;

```

```

    if (strcmp(type, "statreqst") == 0)
        msgtype = STATREQST;
    if (strcmp(type, "tokreqst") == 0)
        msgtype = TOKPREQST;
    if (strcmp(type, "inigrview") == 0)
        msgtype = INITGVIEW;
    if (strcmp(type, "initable") == 0)
        msgtype = INITTABLE;
    if (strcmp(type, "initpool") == 0)
        msgtype = INITPOOL;
    if (strcmp(type, "timeout__") == 0)
        msgtype = TIMEOUT__;
    if (strcmp(type, "startmr") == 0)
        msgtype = STARTTIMR;
    if (strcmp(type, "deltoken") == 0)
        msgtype = DELTTOKEN;
    if (strcmp(type, "extknpool") == 0)
        msgtype = EXTKNPOOL;

    return(msgtype);
} /* end msg_type */

```



```

/*****
get_target - extracts the destination field of an
internal message.
The original message is not disturbed. The result
is stored on a dynamic array 'nbr', that has to
be deallocated before reusing.
Sample call:

char *nbr;
...
nbr = get_target(msg);
...
free(nbr);
*****/

char *get_target(str)
char *str;
{
    char *nbr, *nbrtmp, *tmp;
    link *list;

    tmp = CALLOC(strlen(str) + 1, char);
    strcpy(tmp, str);

    list = str2list(tmp, "\n ");

    if ( getfromlist(list, &nbrtmp, 2) != 2) {
        printf("get_target error\n");
        printf("get_target : string is %s\n", str);
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    nbr = CALLOC(strlen(nbrtmp) + 1, char);
    strcpy(nbr, nbrtmp);

    removelist(list);
    free(tmp);

    return (nbr);
} /* end get_target */

```

```

/*****
get_originator - extracts the originator field from the
header of an external message.
The original message is not disturbed. The result
is stored on a dynamic array 'nbr', that has to
be deallocated before reusing.
Sample call:

char *nbr;
...
nbr = get_originator(msg);
...
free(nbr);
*****/

char *get_originator(str)
char *str;
{
    char *nbr, *nbrtmp, *tmp;
    link *list;

    tmp = CALLOC(strlen(str) + 1, char);
    strcpy(tmp, str);

    list = str2list(tmp, "\n ");

    if ( getfromlist(list, &nbrtmp, 2) != 2) {
        printf("get_originator error\n");
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    nbr = CALLOC(strlen(nbrtmp) + 1, char);
    strcpy(nbr, nbrtmp);

    removelist(list);
    free(tmp);
    return (nbr);

} /* end get_originator */

```


Table of Contents

Socket Utilities	78
createUDP	78
createUN	79
connectUN	79
readn	80
writen	80
readmsg	81
writemsg	82
senmsg	83
recmsg	84


```
/******
```

* SOCKET INTERFACE AUXILIARY FUNCTIONS

*
* The following functions are available to be used:

```
* int createUDP(u_short port);  
* int createUN(char *path);  
* int connectUN(char *server_path);  
* int readmsg(int fd, char **ptr, char *con);  
* int writemsg(int fd, char *ptr, int n);  
* void sendmsg(char *msg, int n, char *IPaddr, u_short port);  
* int recmsg(int fd, char **str,);
```

* Refer to the function header comments for detailed info.
* Other functions in this file are used internally, and should
* not be used directly.

```
*****
```

* Written by: Fernando J. Pires
* Last revision: 18 Feb 1993

```
*****
```

```
/******
```

createUDP - establish an UDP socket for a server

```
*****
```

```
int createUDP(port)  
u_short *port;
```

```
{  
    struct sockaddr_in sin; /* Internet endpoint address */  
    int sockfd; /* socket descriptor */  
    int sinlen;  
  
    bzero((char *)&sin, sizeof(sin)); /* clear address structure */  
    sin.sin_family = AF_INET;  
    sin.sin_addr.s_addr = htonl(INADDR_ANY);  
    sin.sin_port = *port;
```

```
/* Open the socket */  
if ((sockfd = socket(PF_INET, SOCK_DGRAM, 0)) < 0) {  
    printf("\ncreateUDP: can't open internet socket\n");  
    printf("07\n");  
    exit(1);  
}  
  
/* Bind the socket */  
sinlen = sizeof(sin);  
  
if (bind(sockfd, (struct sockaddr *) &sin, sinlen) < 0) {  
    printf("\ncreateUDP: can't bind local address\n");  
    printf("07\n");  
    exit(1);  
}  
  
if (getsockname(sockfd, (struct sockaddr *) &sin, &sinlen) < 0) {  
    printf("\ncreateUDP: can't bind local address\n");  
    printf("07\n");  
    exit(1);  
}  
  
*port = sin.sin_port;  
  
return sockfd;  
} /* end createUDP */
```

```

/*****
    createUN - establish an Unix Domain socket for a server
*****/
int createUN(path)
char* path;

{
    struct sockaddr_un sunx; /* Internet endpoint address */
    int sockfd; /* socket descriptor */
    int sunlen; /* Addr struct length */

    bzero((char*)&sunx, sizeof(sunx)); /* clear address structure */
    sunx.sun_family = AF_UNIX;
    strcpy(sunx.sun_path, path);
    sunlen = strlen(sunx.sun_path) + sizeof(sunx.sun_family);

    /* Open the socket */
    if ((sockfd = socket(AF_UNIX, SOCK_STREAM, 0)) < 0) {
        printf("createUN: can't open unix socket\n");
        printf("\074\n");
        exit(1);
    }

    /* Bind the socket */
    unlink(path); /* in case it was left open by a previous call */
    if (bind(sockfd, (struct sockaddr*)&sunx, sunlen) < 0) {
        printf("createUN: can't bind local path\n");
        printf("\074\n");
        exit(1);
    }

    return sockfd;
} /* end createUN */

```

```

/*****
    connectUN - establish an Unix Domain socket for a client
*****/
int connectUN(server_path)
char* server_path;

{
    struct sockaddr_un sunx; /* Internet endpoint address */
    int sockfd; /* socket descriptor */
    int sunlen; /* Addr struct length */

    bzero((char*)&sunx, sizeof(sunx)); /* clear address structure */
    sunx.sun_family = AF_UNIX;
    strcpy(sunx.sun_path, server_path);
    sunlen = strlen(sunx.sun_path) + sizeof(sunx.sun_family);

    /* Open the socket */
    if ((sockfd = socket(AF_UNIX, SOCK_STREAM, 0)) < 0) {
        printf("connectUN: can't open unix socket. Attempted %s\n", server_path);
        printf("\n ERROR CODE = %d\n", errno);
        printf("\074\n");
        exit(1);
    }

    /* Connect to the server */
    if (connect(sockfd, (struct sockaddr*)&sunx, sunlen) < 0) {
        printf("connectUN: can't connect to unix server\n");
        printf("\074\n");
        exit(1);
    }

    return sockfd;
} /* end connectUN */

```

```

*****
readn - Read "n" bytes from a descriptor
      Use in place of read() when fd is a stream socket
*****
int readn(fd, ptr, nbytes)
register int  fd;
register char *ptr;
register int  nbytes;

{
    int  nleft, nread;

    nleft = nbytes;

    while (nleft > 0) {
        nread = read(fd, ptr, nleft);

        if ( nread < 0)
            return(nread);
        else
            if(nread == 0)
                break;

        nleft -= nread;
        ptr += nread;
    } /* end while */

    return(nbytes - nleft);
} /* end readn */

```

```

*****
writen - Write "n" bytes to a descriptor
      Use in place of write() when fd is a stream socket
*****
int writen(fd, ptr, nbytes)
register int  fd;
register char *ptr;
register int  nbytes;

{
    int  nleft, nwriten;

    nleft = nbytes;

    while (nleft > 0) {
        nwriten = write(fd, ptr, nleft);

        if ( nwriten < 0)
            return(nwriten);

        nleft -= nwriten;
        ptr += nwriten;
    } /* end while */

    return(nbytes - nleft);
} /* end writen */

```

```

/*****

```

```

readmsg - Read a complete message from a descriptor

```

Use in place of read() when fd is a stream socket
The message is assumed to be terminated by 'com'.
The function allocates the necessary space to build
a non-Null terminated string '*str', plus space for
an extra NULL (to be used by the calling function).
The calling function must free the allocated space,
when it is no longer necessary.

Sample call sequence:

```

char *str;

len = readmsg(fd, &str, "#");
str[len] = NULL;

...
free(str);

```

```

*****/

```

```

int readmsg(fd, ptr, com)
register int fd;
register char **ptr;
register char *com;

```

```

{
    int n, rc, maxlen;
    char c, msghead[HEADERSIZE + 1], *tmp;

```

```

    /* get msg size */
    tmp = msghead;

```

```

    do {
        if ((rc = read(fd, &c, 1)) != 1)
            return(0);

```

```

        *tmp++ = c; } while( c != '~');

```

```

    *--tmp = NULL; /* substitute NULL for '~' to end string */

```

```

    if ((maxlen = atoi(msghead)) == 0)
        return(0);

```

```

    /* allocate space for message and extra NULL */
    *ptr = tmp = CALLOC(maxlen + 1, char);

```

```

    /* get message character by character */
    for (n = 1; n <= maxlen; n++) {

```

```

        if ((rc = read(fd, &c, 1)) == 1) {
            *tmp++ = c;

```

```

        } else if (c == 'com')
            break; } /* End of message */

```

```

    } else if (rc == 0) {

```

```

        if (n == 1) {
            free(*ptr);
            return(0); } /* EOF, no data read */

```

```

        else
            break; } /* EOF, data was read */
    /* Note: the calling function has to free the allocated space */

```

```

    } else {
        free(*ptr);
        return(-1); } /* error */

```

```

    } /* end for */

```

```

    return(n);

```

```

} /* end readmsg */

```

```

/*****

```

```

writemsg - Writes a complete message 'ptr' of size 'n' to a
file descriptor 'fd'. It appends an header that
contains the size of the original message, plus a '-'
as a separator. This header is to be processed by
readmsg().

```

```

It returns the number of characters from
the original message that were actually transmitted.
The original message is not changed.

```

Sample call:

```

n = writemsg(fd, str, strlen(str));

```

```

*****/

```

```

int writemsg(fd, ptr, n)
register int fd;
register char *ptr;
register int n;

```

```

{
    int nwrite, len;
    char header[HEADERSIZE + 1], *msg;

```

```

    if (n > (len = strlen(ptr)))
        n = len;

```

```

    sprintf(header, "%d-", n);
    msg = CALLOC(n+strlen(header) + 1, char);
    strcpy(msg, header);
    strcat(msg, ptr, n);
    nwrite = write(fd, msg, strlen(msg));
    free(msg);

```

```

    return(nwrite-strlen(header));

```

```

} /* end writemsg */

```

```

/*****
semmsg - sends an external message 'msg' of size 'n'
to the specified IP destination <IPAddr, port>.
An header with the value of the size of the
message, and a '.' as separator, is appended
to the message. T. is header is to be processed
by recmsg().
The original message is not disturbed.

Sample call:
    semmsg(msg, n, IPAddr, port);
*****/
void semmsg(msg, n, IPAddr, port)
char *msg;
int n;
char *IPAddr;
u_short port;
{
    struct sockaddr_in target_addr;
    struct hostent *phe;
    struct iovec iov[2];
    int sockfd, len;
    u_short output;
    char header[HEADER_SIZE];

    /* get the target UDP socket description */
    bzero((char *)&target_addr, sizeof(target_addr)); /* clear address structure */
    target_addr.sin_family = AF_INET;
    target_addr.sin_port = port;

    if ((target_addr.sin_addr.s_addr = inet_addr(IPAddr)) == INADDR_NONE)
    {
        if (phe = gethostbyname(IPAddr))
            bcopy(phe->h_addr, (char *)&target_addr.sin_addr, phe->h_length);
        else {
            printf("semmsg: can't get IP host entry\n", IPAddr);
            exit(1);
        }
    }
}

```

```

/* set a connection to the destination */
output = htons(0);
sockfd = createUDP(&output); /* request an arbitrary socket */
connect(sockfd, (struct sockaddr *)&target_addr, sizeof(target_addr));

/* assemble a scattered message including the msg size */
if (n > (len = strlen(msg)))
    .. = len;

iov[0].iov_base = header;
sprintf(header, "%d.", n);
iov[0].iov_len = strlen(header);
iov[1].iov_base = msg;
iov[1].iov_len = n;

/* send message */
if (writev(sockfd, &iov[0], 2) != (strlen(header) + n)) {
    printf("semmsg: write error on socket\n");
    exit(1);
}
close(sockfd);
} /* end semmsg */

```

```

/*****
recmsg - reads an external message 'msg' at the
specified IP socket.
The message is atomically received, and is stripped
of the header (as created by sendmsg()).
The function allocates the necessary space to build
a non-Null terminated string *str, plus space for
an extra NULL (to be used by the calling function).
The calling function must free the allocated space,
when it is no longer necessary.
The function returns the message size, and -1 if
an invalid message is received.
Sample call sequence:
char *str;
len = recmsg(fd, &str);
str[len] = NULL;
...
free(str);
*****/

```

```

int recmsg(fd, str)
register int fd;
register char **str;
{
    char *msghead, *tmp, *msgbuf, buf[HEADERSIZE + 1];
    int msglen, recrlen, mlen;

    /* retrieve the header (the message is not removed) */
    if ((msglen = recv(fd, buf, HEADERSIZE, MSG_PEEK)) < 0) {
        printf("recmsg: header error\n");
        exit(1);
    }

    buf[HEADERSIZE] = NULL;
    msghead = strtok(buf, "~");

    if ((msglen = atoi(msghead)) == 0)
        return(0);

```

```

    /* allocate space for entire message */
    msgbuf = CALLOC(msglen + HEADERSIZE + 1, char);

    /* get entire message */
    if ((recrlen = recv(fd, msgbuf, msglen + HEADERSIZE, 0)) < 0) {
        printf("recmsg: message error\n");
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        free(msgbuf);
        return(-1);
    }

    msgbuf[recrlen] = NULL;

    /* extract message info and discard header */
    tmp = strtok(msgbuf, "~");
    tmp = strtok(NULL, "~");
    mlen = strlen(tmp);
    *str = CALLOC(mlen + 1, char);
    strcpy(*str, tmp);
    free(msgbuf);

    return(mlen);
} /* end recmsg */

```

**A SIMPLE APPLICATION
and
MAIN PROCESS**

Simple Application	86
Main Process	88


```

/*****
* GROUP MEMBERSHIP PROTOCOL - SIMPLE APPLICATION
*
* This is an example application, that creates an instance of the
* membership protocol, and receives from it the most current group
* view, when this changes.
* A mechanism for requesting the element shut-down is also
* provided. This is accomplished with the following call:
*
* kill(0, SIGALRM);
*
* The application receives notice that the element has departed
* by catching this signal (sent by mainproc) at function killelmt()
*****/
* Written by: Fernando J. Pire
* Last revision: 9 Mar 1993
*****/

```

```

#include "grp.h"
#include "socutil.c"

```

```

void killelmt();

```

```

void main(argc,argv)

```

```

int  argc;
char *argv[];

```

```

{
    int  apsocun, newsoc, childpid, clen, msglen;
    char *apath, *groupname, *sitelist, *msg;
    struct sockaddr_un caller_addr;

```

```

/*****
    DETERMINE INPUT ARGUMENTS (command line)
*****/
switch(argc){

```

```

    case 1:
        groupname = "group0";
        sitelist = (char*) NULL;
        break;

    case 2:
        groupname = argv[1];
        sitelist = (char*) NULL;
        break;

    case 3:
        groupname = argv[1];
        sitelist = argv[2];
        break;

    default:
        printf("usage: simpleapp [groupname [sitelist ]Nt");
        exit(-1);
    } /* end switch */

/*****
    OPEN LOCAL SOCKET (where group views are received)
*****/
apath = UNIXSTR_TMPL;
mktemp(apath);
apsocun = createUN(apath);
listen(apsocun, 5);

/*****
    ESTABLISH A SIGNAL HANDLER TO INTERCEPT
    THE ELEMENT FAILURE SIGNAL FROM mainproc
*****/
signal(SIGALRM, killelmt);

```

```

*****
EXECUTE GMP's MAIN PROCESS
*****
if ((childpid = fork0) == -1)
    printf("Can't fork\n");

else
    if (childpid == 0) { /* child process */
        execvp("mainproc", "mainproc", ap1path, groupname, stielist, (char**)NULL);
        printf("Error executing mainproc\n");
        exit(1); }

*****
EXECUTE LOOP TO RECEIVE UPDATED GROUP VIEWS
*****
while ((newsoc = accept(ap1socun, (struct sockaddr*) &caller_addr, &clen)) >= 0) {
    if((msglen = recvmsg(newsoc, &msg)) < 0) {
        printf("APPLICATION: read error\n");
        break; }

    msg[msglen] = NULL; /* turn message into string */
    printf("#####\n");
    printf("#####\n");
    printf("APPLICATION: received group view => #%s\n", msg);
    printf("#####\n");
    printf("#####\n");
    free(msg);
    close(newsoc);
} /* end while */

printf("APPLICATION: accept error\n");

/* Send signal to mainproc and to itself, requesting element shut-down */
kill(0, SIGALRM);

} /* end simpleapp */
}

*****
SIGNAL HANDLER THAT CATCHES ELEMENT DEPARTURE
*****
void killelmt()
{
    printf("APPLICATION: mainproc has returned\n");

    /* Terminate all running processes */
    sleep(2); /* Allow time for mainproc to shut-down */
    kill(0, SICKILL);
} /* end killelmt */

```

```

/*****
* GROUP MEMBERSHIP PROTOCOL - MAIN PROCESS
*
* This program is executed by the application, and spawns
* complete implementation of an element running the Membership Protocol
* This program waits for a child to cease execution (meanin
* that the element has or is to cease existence, and then releases all
* resources to the operating system.
* It also catches signals from the application requesting the
* element shut-down.
*
* Example of code used by the application to run this program:
*
* //create unix socket where GroupViews are to be received/* char
* socpath;
* int socfd;
*
* socpath = UNIXSTR_TMPL; // Default path template //
* mktemp(socpath); // Get unique file name //
* socfd = createUN(socpath); // Create unix socket //
* listen(socfd,5);
*
* //fork and execute mainproc//
* if ( (childpid = fork0) == -1 )
* printf("Can't fork\n");
* else if (childpid == 0){ // child process //
* execvp("mainproc", "mainproc", socpath, grouppathname, sitelist,
* (char*)NULL);
* printf("Error executing front\n");
* exit(1);;
*
* Notes: sitelist is a string with host names, separated by
* '=' characters. Example: "sun2=sun10=aditya=taurus"
* This list can be empty, in which case the local host
* Groupname is optional. If no argument is provided
* it defaults to 'group0'.
*
*****/

```

```

*****
* Written by: Fernando J. Pires
* Last revision: 9 Mar 1993
*
*****
#include "grp.h"
#include "socutil.c"

void killelmt0;

/* these variables are common to 'mainproc' and 'killelmt' */
char *fpath, *bpath, *smonpath, *streppath, *timerpath, *joinpath,
*intmbpath, *agppath, *comppath, *gvmpath, *smpath,
*tmppath, *apfpath, *grouppathname;

int main(argc,argv)
int argc;
char *argv[];
{
    int fsocudp, fsocun, bsocudp, bsocun, smonfd, strepfd, timerfd,
    joinpfd, intmbfd, agrpfd, compfd, gvmpfd, smpfd, tmppfd,
    commpid, frontpid, backpid, gvmppid, smpid, tmppid, timerpid,
    streppid, smonpid, intmbpid, agreepid, joinpid, returnpid;
    char sfdup[MAXFD], stdup[MAXFD], sfun[MAXFD], sbun[MAXFD], ssmon[MAXFD],
    ssrep[MAXFD], stimer[MAXFD], sjoinp[MAXFD], sintmb[MAXFD],
    sagrp[MAXFD], scomp[MAXFD], sgvm[MAXFD], sstm[MAXFD], stpm[MAXFD],
    sfport[MAXPORT], sbport[MAXPORT], my_name[MAXHOSTNAME + 1],
    *ip_addr, my_addr[MAXLMTSIZE + 1], *groupname, *sitelist;

    u_short front, bport;
    struct hostent *hptr;

```



```

/* open BACK port UDP socket ( an Internet Datagram Socket) */
bport = htons(0);
bsocudp = createUDP(&bport);
sprintf(sbuudp, "%d", bsocudp);
printf("BACK: internet port => %d\n", ntohs(bport));
sprintf(sbport, "%d", bport);
printf("\n");

/* open a BACK port Unix Domain Stream socket */
bpath = UNIXSTR_TMPL; /* Default path template */
mktemp(bpath); /* Get unique file name */
bsocun = createUN(bpath);
listen(bsocun, 5);
sprintf(sbuun, "%d", bsocun);
printf("BACK: unix socket path => %s\n", bpath);
printf("\n");

/* open a STATUS MONITOR Unix Domain Stream socket */
simonpath = UNIXSTR_TMPL; /* Default path template */
mktemp(simonpath); /* Get unique file name */
simonfd = createUN(simonpath);
listen(simonfd, 5);
sprintf(ssmon, "%d", simonfd);
printf("STATUS MONITOR: unix socket path => %s\n", simonpath);
printf("\n");

/* open a STATUS REPORTER Unix Domain Stream socket */
sreppath = UNIXSTR_TMPL; /* Default path template */
mktemp(sreppath); /* Get unique file name */
srepdf = createUN(sreppath);
listen(srepdf, 5);
sprintf(ssrep, "%d", srepdf);
printf("STATUS REPORTER: unix socket path => %s\n", sreppath);
printf("\n");

/* open a TIMER Unix Domain Stream socket */
timerpath = UNIXSTR_TMPL; /* Default path template */
mktemp(timerpath); /* Get unique file name */
timerfd = createUN(timerpath);

```



```

(char*) NULL);
printf("Error executing join\n");
printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
exit(1); }

/* execute INTEGRATE MEMBER process */
if ((inimbrpid = fork()) == -1)
    printf("Can't fork\n");
else
    if (inimbrpid == 0) { /* child process */
        execlp("inimbr", "inimbr", simnbr, my_addr, gvmppath, stmpath,
            tmpath, bpath, (char*) NULL);
        printf("Error executing inimbr\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(1); }

/* execute AGREEMENT PROCESSOR process */
if ((agrecpid = fork()) == -1)
    printf("Can't fork\n");
else
    if (agrecpid == 0) { /* child process */
        execlp("agree", "agree", sgrp, my_addr, stmpath, gvmppath, tmpath,
            joinppath, comppath, fpath, (char*) NULL);
        printf("Error executing agrp\n");
        printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(1); }

/* close all open files */
close(fsocudp);
close(hsocudp);
close(fsocun);
close(hsocun);
close(simonid);
close(strepid);
close(timerfd);
close(joinpid);
close(inimbrfd);
close(agrpfd);
close(compid);

close(gvmfd);
close(smfd);
close(tpmfd);

/* wait until one process exits */
returnpid = wait((int*) NULL);

printf("\07MAIN PROCESS: ");

if (returnpid == commppid)
    printf("Commit");
if (returnpid == fronpid)
    printf("FRONT");
if (returnpid == backpid)
    printf("BACK");
if (returnpid == gvmppid)
    printf("GVM");
if (returnpid == simpid)
    printf("STM");
if (returnpid == tpmppid)
    printf("TPM");
if (returnpid == timerpid)
    printf("Timer");
if (returnpid == streppid)
    printf("Status Reporter");
if (returnpid == simonppid)
    printf("Status Monitor");
if (returnpid == inimbrpid)
    printf("Integrate Mbr");
if (returnpid == agreepid)
    printf("Agree");

printf(" has returned (pid = %d)\n", returnpid);
printf("XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");

sleep(120);

printf("MAIN PROCESS: ");

```



```

if (returnpid == commpid)
    printf("Committ");
if (returnpid == frontpid)
    printf("FRONT");
if (returnpid == backpid)
    printf("BACK");
if (returnpid == gvmppid)
    printf("GVM");
if (returnpid == sumpid)
    printf("STM");
if (returnpid == ipmpid)
    printf("TPM");
if (returnpid == timerpid)
    printf("Timer");
if (returnpid == streppid)
    printf("Status Reporter");
if (returnpid == simonpid)
    printf("Status Monitor");
if (returnpid == intmbpid)
    printf("Integrate Mbr");
if (returnpid == agreepid)
    printf("Agree");

```

printf(" has returned (pid = %d)\n", returnpid);

/* Remove Unix Domain socket links */

```

unlink(fpath);
unlink(bpath);
unlink(simonpath);
unlink(streppath);
unlink(integratepath);
unlink(joinppath);
unlink(intmbpath);
unlink(agreepath);
unlink(compopath);
unlink(gvmpath);
unlink(sumpath);
unlink(ipmpath);
unlink(aplpath);

```

```

/* Signal the application that the element has ceased existence */
kill(0, SIGALRM);

} /* end mainproc */

/******
   SIGNAL HANDLER THAT CATCHES ELEMENT DEPARTURE
   *****/
void killelmt()

{
    printf("MAINPROC: application has requested shut-down\n");

    /* Remove Unix Domain socket links */
    unlink(fpath);
    unlink(bpath);
    unlink(simonpath);
    unlink(streppath);
    unlink(timerpath);
    unlink(joinppath);
    unlink(intmbpath);
    unlink(agreepath);
    unlink(compopath);
    unlink(gvmpath);
    unlink(sumpath);
    unlink(ipmpath);
    unlink(aplpath);
    remove(grouppathname);
    free(grouppathname);

} /* end killelmt */

```

FIFO CHANNEL

FIFO Channel Process Dependencies	96
Front Process Specification	97
Back Process Specification	98
Front	99
Back	102




Figure A1 FIFO Channel - Process Dependencies

FIFO Channel - *FRONT* Process

```
1  Wait for a channel ready to read
2  if (external channel ready)
3      if (Status_Query)
4          send Status_Query to MonitorProcess
5      else if (JoinRequest)
6          send JoinRequest to JoinProcessor
7      else if (InitialParameters)
8          send InitialParameters to JoinProcessor
9      else if (TokenAck)
10         if (Received_Serial_Number = Expected_serial_number)
11             remove Head_of_Queue
12             decrement Queue_Counter
13         end
14     end
15 else /* internal channel ready */
16     if (Token)
17         change Token to external format /* add external header */
18         insert Token in queue
19         increment Serial_Number
20         increment Queue_Counter
21     else if (TokenPool)
22         discard all messages in queue
23         change TokenPool to external format /* add external header */
24         insert TokenPool in queue
25         increment Serial_Number
26         increment Queue_Counter
27     else if (StatusReport)
28         update cwnbr
29         send StatusReport to cwnbr
30     end
31 end
32 if (Queue_Counter > 0)
33     send Head_of_Queue to cwnbr
34     set Expected_serial_number = Head_of_Queue_serial_number
35 end
```

Figure A2 FIFO Channel - Front Process

FIFO Channel - *BACK* process

```
1  Wait for a channel ready to read
2  if (internal channel ready)
3      if (Status_Query)
4          update acwnbr
5          send Status_Query
6      else if (Initial_Parameters)
7          update acwnbr
8          send Initial_Parameters
9      else if (Join_Request)
10         send Join_Request
11     end
12 else /* external channel ready */
13     if (message originator = acwnbr)
14         if (Status_Report)
15             send Status_Report to MONITOR_PROCESS
16         else if (Token)
17             if (Serial_Number = Expected_Serial_Number - 1)
18                 send Token_Ack /* to acwnbr */
19             end
20             if (Serial_Number = Expected_Serial_Number)
21                 send Token to AgreeProcessor
22                 send Token_Ack /* to acwnbr */
23                 increment Expected_Serial_Number
24             end /* out of order messages are discarded */
25         else if (Token_Pool) /* Token_Pool is always accepted */
26             send Token_Pool to AgreeProcessor
27             send Token_Ack /* to acwnbr */
28             set Expected_Serial_Number = Serial_Number + 1
29         end
30     end
31 end
```

Figure A3 FIFO Channel - Back Process


```
close(newsoc);
```

```

default:
    printf("FRONT: invalid message type %d\n", msgtype);
    printf("07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(1);

} /* end switch */

free(msg);

} /* end else if (FD_ISSET(sockudp, &fdread) */

if (queue_counter > 0) {
    get_queue_head(head, &tmpmsg);

    tmpmsg = CALLOC (sizeof(tmpmsg) + 1, clear);
    strcpy(tmpmsg,tmpmsg);

    send_msg_back(tmpmsg, cwnbr);
    expectsnbr = get_sr_nbr(tmpmsg);

    free(tmpmsg);

} /* if (queue_counter > 0) */

} /* end while TRUE */

} /* end FRONT */

```



```
case EXTKNPOOL:
    send_msg_in(msg, agtp);
    send_ack(msg, my_addr, acwnbr);

    lastsnbr = get_sr_nbr(msg);
    expectsnbr = lastsnbr + 1;
    break;

default:
    printf("BACK: received invalid message type\n");
    printf("\0/XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
    exit(1);

    } /* end switch */

    } /* end if (strcmp(acwnbr,originator) */

    free(originator);
    free(msg);

    } /* end if (FD_ISSET(sockudp, &fdread)) */

    } /* end while */

    } /* end back */
```

MONITOR PROCESS

Monitor Process Dependencies	106
Status Reporter Process Specification	106
Status Monitor	107
Satus Reporter	110
Timer	113

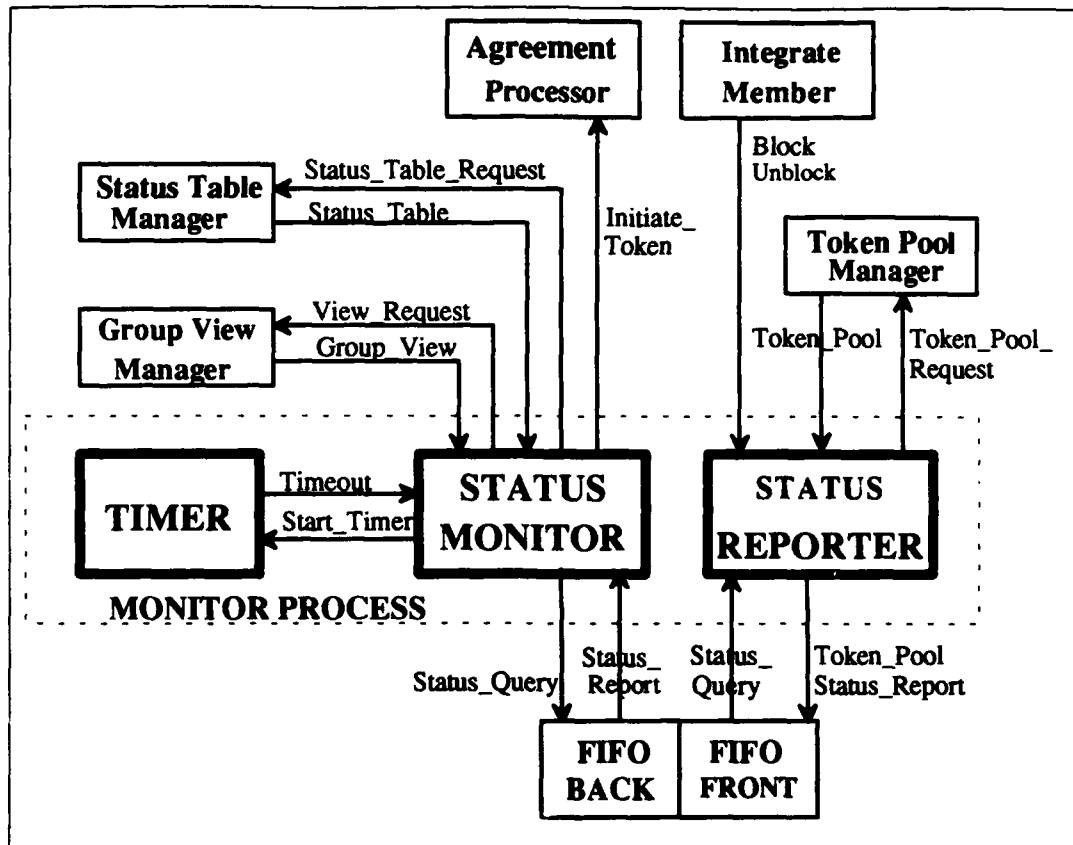


Figure A4 Monitor Process - Internal Structure and Dependencies

ReportStatus process at p_i

```

1  if (not blocked by IntegrateMember)
2    if (querying member  $\in GV_{p_i}$  or has joinagree status)
3       $p_{mon}$  = querying member
4      send status to  $p_{mon}$ 
5      if (previous querying member =  $p_{mon}$ )
6        send TokenPool( $p_i$ ) to  $p_{mon}$ 
7      end
8    end
9  end
end ReportStatus

```

Figure A5 Reporting of Status


```

close(timfd);

/* Accept query expiration msg from timer */
clen = sizeof(caller_addr);

if ((newsoc = accept(soc, (struct sockaddr*)&caller_addr, &clen)) < 0) {
    printf("\nStatusMonitor: accept from timer failed.\n");
    printf("\07$");
    exit(-1);
}

if ((msglen = readmsg(newsoc, &buf, "#")) < 0) {
    printf("\nStatusMonitor: timer read failed.\n");
    printf("\07$");
    exit(-1);
}

buf[msglen] = NULL;
close(newsoc);

msgfromtimer = str2list(buf, "\n#");

if (getfromlist(msgfromtimer, &timermstype, 2) == 0) {
    printf("\nStatusMonitor: timer msg parsing failed.\n");
    printf("\07$");
    exit(-1);
}

if (strcmp(timermstype, "lqry") != 0) {
    printf("\nStatusMonitor: timer func failed.\n");
    printf("\07$");
    exit(-1);
}

/* end if ( strcmp(timermstype, "lqry") != 0) */
remove(list(msgfromtimer);
free(buf);
break;

default:
    printf("\nStatusMonitor: received invalid msg type.\n");
    printf("\07$");
    exit(-1);

} /* end switch */
} /* end while TRUE */
} /* end main */

```

```

close(timfd);

/* Accept query expiration msg from timer */
clen = sizeof(caller_addr);

if ((newsoc = accept(soc, (struct sockaddr*)&caller_addr, &clen)) < 0) {
    printf("\nStatusMonitor: accept from timer failed.\n");
    printf("\07$");
    exit(-1);
}

if ((msglen = readmsg(newsoc, &buf, "#")) < 0) {
    printf("\nStatusMonitor: timer read failed.\n");
    printf("\07$");
    exit(-1);
}

buf[msglen] = NULL;
close(newsoc);

msgfromtimer = str2list(buf, "\n#");

if (getfromlist(msgfromtimer, &timermstype, 2) == 0) {
    printf("\nStatusMonitor: timer msg parsing failed.\n");
    printf("\07$");
    exit(-1);
}

/*
 * The following discards a trap msg from timer that is delayed.
 * A two-level if is used instead of a while because there is
 * exactly one trap to be discarded if any.
 */
if (strcmp(timermstype, "lqry") != 0) {
    remove(list(msgfromtimer);
    free(buf);
    clen = sizeof(caller_addr);
}

```



```

/*****
* StatusReporter.
* Algorithm: Fig. 4 of the gmp paper
* Description: Uses one Unix domain socket. Receives FRONT and token pool manager
*             socket names along with addr of the member it belongs to.
* Written by: Shridhar Shukla
*             David Pezdirtz
*
* Date: 27 Jul 1993
*
*****/

#include "gmputil.c"

main(argc,argv)
int argc;
char *argv[];
{
    int soc, newsoc, clen, msglen, fromfd, fromfdt;
    char *frontsoc, *tpmsoc, *myaddr, mon[MAXLMTSIZE], prev_mon[MAXLMTSIZE],
    *querydata, *buf, *target, *originator, itprec[NODATAMSG+1],
    *msg, *pool;
    link *query, *query_components, *list;
    struct sockaddr_un caller_addr;

    printf("\nStatusReporter Process: start execution\n");

    /* Get socket file descriptors from command line argument */
    if (argc == 5){
        soc = atoi(argv[1]);
        myaddr = argv[2];
        tpmsoc = argv[3];
        frontsoc = argv[4];
    }
    else{

```

```

        printf("\nUsage: StatusReporter soc myaddr tpmsoc frontsoc\n");
        printf("\n07$SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    strcpy(prev_mon, myaddr); /* initially, one monitors self */

    /*****
    * Wait for a query to appear and respond to it.
    *****/

    while(TRUE){
        clen = sizeof(caller_addr);

        /* Accept status query from front */
        clen = sizeof(caller_addr);

        if ((newsoc = accept(soc, (struct sockaddr*)&caller_addr, &clen)) < 0){
            printf("\nStatusReporter unix: accept error\n");
            printf("\n07$SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
            exit(-1); }

        /* Read message */
        if ((msglen=readmsg(newsoc,&buf,&#")) < 0) {
            printf("\nStatusReporter: unix PORT read error\n");
            printf("\n07$SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
            exit(-1); }

        buf[msglen]=NULL;

        query = str2list(buf, "\n#");

        if ( getfromlist(query, &querydata, 2) == 0){
            printf("\nStatusReporter: query parsing failed\n");
            printf("\n07$SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
            exit(-1); }

        query_components = str2list(querydata, "");

```

```

if (strcmp(mon, prev_mon) != 0) { /* if there is a new monitor */
    strcpy(prev_mon, mon); /* update the previous monitor */

    /* get the token pool */
    strcpy(tkreq, "tokreqst#");
    msglen = strlen(tkreq);
    tpmfd = connectUN(tpmsoc);

    if ( writemsg(tpmfd, tkreq, msglen) < msglen ) {
        printf("\nStatusReporter: token pool request failed\n");
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    if ( (msglen=readmsg(tpmfd,&buf,"#")) < 0 ) {
        printf("\nStatusReporter: token pool read failed\n");
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    buf[msglen] = NULL;
    close(tpmfd);

    /* create the external token pool message */
    list = str2list(buf, "\n#");

    if ( getfromlist(list, &pool, 2) == 0 ) {
        printf("\nStatusReporter: token pool parsing failed\n");
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    msg = CALLOC(9 + strlen(myaddr) + strlen(pool) + 3 + 1, char);
    strcpy(msg, "extknpool\n");
    strcat(msg, myaddr);
    strcat(msg, "\n");
    strcat(msg, pool);
    strcat(msg, "#");

    /* send the token pool to monitor via front */
    msglen = strlen(msg);
    frontfd = connectUN(frontsoc); /* open connection */

```

```

if ( getfromlist(query_components, &target, 1) == 0 ) {
    printf("\nStatusReporter: query component parsing failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

if (strcmp(myaddr, target) != 0) {
    printf("\nStatusReporter: target of query is not me\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

if ( getfromlist(query_components, &originator, 2) == 0 ) {
    printf("\nStatusReporter: query component parsing failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

strcpy(mon, originator);
removelst(query);
removelst(query_components);

/* assemble a status report message in the same place as the query received */
strcpy(buf, "statusrpt"); /* msg type */
strcat(buf, "\n"); /* separator */
strcat(buf, mon); /* target of report */
strcat(buf, "\n"); /* separator */
strcat(buf, myaddr); /* originator of report */
strcat(buf, "#"); /* end of message delimiter */

/* send the status report out */
msglen = strlen(buf);
frontfd = connectUN(frontsoc); /* open connection */

if ( writemsg(frontfd, buf, msglen) < msglen ) { /* write all the bytes */
    printf("\nStatusReporter: report send failed\n");
    printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

close(frontfd);

free(buf); /* storage acquired on read of query returned */

```

```
if ( writemsg(frontid, msg, msglen) < msglen ) {  
    printf("StatusReporter: token pool send failed.\n");  
    printf("\073XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");  
    exit(-1);  
    close(frontid);  
    remove(list);  
    free(buf);  
    free(msg);  
    } /* end if (strcmp(mon, prev_mon) != 0) */  
    close(newsoc);  
    } /* end while(TRUE) */  
    } /* end StatusReporter */
```



```

*****
GetNextTqry0: The count down loop in Timer.c does not attempt to read the
socket while it is counting down from the value returned by
this function.
*****
int GetNextTqry0
{
    int    next;

    next = TQRY_INTERVAL;
    return(next);
} /* end GetNextTqry */

```

```

*****
GetNextTpad0: The count down loop in Timer.c attempts to reads the
socket while it is counting down from the value returned by
this function. If a msg is read, the down counter is initialized
to the value returned by GetNextTqry0.
*****
int GetNextTpad0
{
    int    next;

    next = TPAD_INTERVAL;
    return(next);
} /* end GetNextTpad */

```

AGREEMENT PROCESSOR

Agreement Processor - Process Dependencies	117
AgreeProcessor Specification	118
Determination of Token Originator's Failure	119
Processing Agree Tokens	120
Agree.c	121

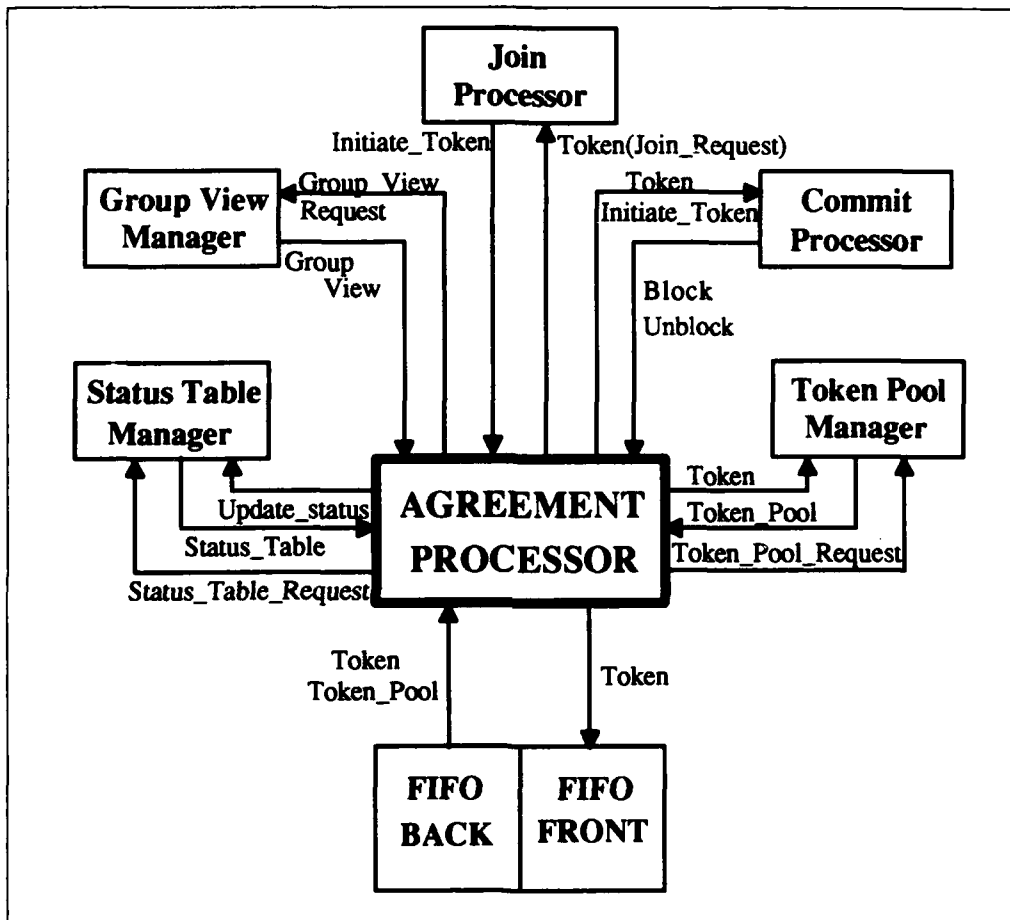


Figure A6 Agreement Processor - Process Dependencies


```

AgreeProcessor for  $agree_{p_j}(p_k)$  at  $p_i$ 
1  if (not blocked by CommitProcessor)
2    if (initiate agreement message received) /*  $p_i = p_j$  */
3      add  $agree_{p_j}(p_k)$  to TokenPool( $p_i$ )
4       $ST_{p_i}(p_k) \leftarrow joinagreed$  or failagreed
5      send  $agree_{p_j}(p_k)$  to cwnbr( $p_i$ )
6      send acknowledgment to calling process
7    else /* a token or external token pool is received */
8      if (ExtTokenPool)
9        for  $\forall tokens \in ExtTokenPool$ 
10         if ( $token \in TokenPool(p_i)$ )
11           if (originator failed)
12             ProcessToken
13           end
14         else /* token not in TokenPool */
15           if (received for the first time)
16             ProcessToken
17           end
18         end
19       end
20     else /* a token was received */
21       if (received for the first time)
22         ProcessToken
23       end
24     end
25   end
26 end

```

Figure A7 Agreement Processor

```

LostAgreeToken
1  if (joinagree)
2    if (rank( $p_j$ ) > rank( $p_i$ ))
3      return true
4    else
5      return false
6    end
7  end

8  if (failagree)
9    if (RelativeRank( $p_k$ ,  $p_l$ ) > RelativeRank( $p_j$ ,  $p_i$ ))
10     return true
11   else
12     return false
13   end
14 end

```

Figure A8 Determination of Token Originator's Failure

```

ProcessToken
1  if (joinreqst)
2      send token to JoinProcessor
3  elseif (commit)
4      send token to ComitProcessor
5  elseif (agree)
6      if ( $(p_i \neq p_j) \ \&\& \ (\textit{agree token} \notin \textit{TokenPool}(p_i))$ )
7          add  $\textit{agree}_{p_j}(p_k)$  to TokenPool( $p_i$ )
8           $ST_{p_i}(p_k) \leftarrow \textit{FailAgreed}$  or JoinAgreed
9          send  $\textit{agree}_{p_j}(p_k)$  to cwnbr( $p_i$ )
10     else
11          $p_j$ 
12         if ( $(p_i = p_j) \ \parallel \ (\forall p_l \mid p_l \rightarrow p_i, p_l \in ST_{p_i})$ )
13             compute rank  $\forall p_l \in ST_{p_i}$  with Agreed status
14             if rank( $p_k$ ) = smallest
15                 send initiate_comit to ComitProcessor
16             else
17                  $ST_{p_i}(p_k) \leftarrow \textit{joinpendg}$  or failpendg
18             end
19         end
20     end

```

Figure A9 Processing Agree Tokens

```
*****
/*****
* Agreement:
* Description: Refer to Fernando's thesis, Chapter 3
* Written by: Shridhar Shukla
*              David Pezdirtz
*
* Date: 14 Dec 1993
*****
```

```
#include "gmputil.c"
```

```
void IntegrateToken0;
void ProcessTokenPool0;
void ProcessToken0;
void ExecuteAgreement0;
int OriginatorFailed0;
```

```
main(argc,argv)
int  argc;
char *argv[
```

```
{
    int soc, newsoc, clen, msglen, msgtype,
        frontid;
    char *myaddr, *sinsoc, *gvmisoc, *ipmsoc, *jpsoc, *comsoc,
        *fronsoc, *buf, token[TOKENMSGLEN+1], *tlen,
        c[1], *initkmsg;
```

char *gv, *st;

```
struct sockaddr_un caller_addr;
link *msglist, *tlink;
```

```
printf("\nAgreement Process: start execution.\n");
```

buf[msglen] = NULL;

```
msgtype = in_msg_type(buf);
```

```

/* create agree token if initiate token: type not checked yet. */
if ( msgtype == INITTOKEN ) {

    /* assemble token */
    msglist = str2list(buf, "\n#");

    if ( getfromlist(msglist, &initkmsg, 2) == 0 ) {
        printf("\nAgreement: parsing failed for initiate token msg.\n");
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    strcpy(token, "tokenok\n");
    strcat(token, initkmsg);
    strcat(token, " ");
    strcat(token, myaddr);
    strcat(token, "\n#");
    removelist(msglist);

    IntegrateToken(token, tpmsock, smsoc); /* update local state */

    /* send token to front */
    msglen = strlen(token);
    frontid = connectUN(frontsoc);

    if ( writemsg(frontid, token, msglen) < msglen ) {
        printf("\nAgreement: token send to front failed.\n");
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    close(frontid);

    /* send NULL ack to process requesting initiate token msg */
    c[0] = NULL;
    msglen = strlen(c);

    if ( writemsg(newsoc, c, msglen) < msglen ) {
        printf("\nAgreement: ack send to token initiator failed.\n");
        printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

```

```

} /* end if INITTOKEN */
else { /* a token is received */
    switch (msgtype) {
        case EXTKNPOOL:

            ProcessTokenPool(buf, myaddr, tpmsock, smsoc, gvmsock, ipsoc, comsoc, frontsoc);
            break;

        case TOKENTOKEN:
            tknl = str2list(buf, "\n#"); /* remove the msg header */

            if ( getfromlist(tknl, &tkn, 2) == 0 ) {
                printf("\nAgreement: parsing failed for a token msg.\n");
                printf("\n07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
                exit(-1); }

            /* get local group view */
            gv = GetGroupView(gvmsock);

            /* get status table */
            st = GetStatusTable(smsoc);

            if (first_time(tkn, gv, st)) {
                ProcessToken(tkn, myaddr, tpmsock, smsoc, gvmsock, ipsoc, comsoc, frontsoc);
                free(gv);
                free(st);
            }

            removelist(tknl);

        } /* end switch */

    } /* end else */

    close(newsoc);
    free(buf);

    } /* end while (infinite loop) */
} /* end main */

```



```

/*****
ProcessToken:
*****/
void ProcessToken(tkn, myaddr, ipmsoc, smsg, smsglen, gvmsock, ipmsoc, comsoc, frontsoc)
char *tkn, *myaddr, *ipmsoc, *smsg, *gvmsock, *ipmsoc, *comsoc, *frontsoc;
{
    link *link;
    char *subject, *stbuf, *origin, *gv, *sa, ltkn[TOKENMSGLEN + 1],
    tkmsg[TOKENMSGLEN + 1], request[NODATAMSG + 1],
    status[STSTYPELEN + 1], *msg;
    int tkntype, msglen, comid, smsglen, smid, jpid, flag;

    strcpy(link, tkn);
    tkn = str2list(link, " ");

    if (getfromlist(link, &subject, 2) == 0) {
        printf("\nProcessToken: parsing failed for token subject\n");
        printf("\n*****\n");
        exit(-1);
    }

    if (getfromlist(link, &origin, 3) == 0) {
        printf("\nProcessToken: parsing failed for token originator\n");
        printf("\n*****\n");
        exit(-1);
    }

    tkntype = GetTokenType(tkn);

    switch(tkntype) {
        case FAILCOMMIT: /* processing same for comit tokens */
        case JOINCOMMIT:
            strcpy(linkmsg, "tokenok\n"); /* send token to commit process */
            strcat(linkmsg, tkn);
            strcat(linkmsg, "#");

            msglen = strlen(linkmsg);
            comid = connectUN(comsoc);

            if (writemsg(comid, linkmsg, msglen) < msglen) {
                printf("\nProcessToken: tkn send to commit process failed\n");
                printf("\n*****\n");
                exit(-1);
            }

            /* wait for response back from commit processor */

            if ((msglen=readmsg(comid, &msg, "#")) < 0) {
                printf("\nExecuteAgreement: fail on response from commit\n");
                printf("\n*****\n");
                exit(-1);
            }

            msg[msglen] = NULL;

            /* don't check message as only possible is completion flag from commit */
            free(msg);

            close(comid);
            break;

        case JOINRQSTT:
            if (InGroup(gvmsock, subject) == -1) { /* subject is NOT in gv */
                /* get the status table */
                stbuf = GetStatusTable(smsg);

                /* and subject is NOT in st */
                if (InStatusTable(stbuf, subject, (char *) NULL) == FALSE) {
                    /* send token to join process */
                    strcpy(linkmsg, "tokenok\n");
                    strcat(linkmsg, tkn);
                    strcat(linkmsg, "#");

                    msglen = strlen(linkmsg);
                    jpid = connectUN(ipmsoc);

```



```

        strcpy(status, "joinagree") == 0)
        strcpy(updstsmg, "joinpendg");
    else if (strcpy(updstsmg, "failpendg");
    else {
        printf("\tExecute Agreement: incorrect status %s\n", status);
        printf("\t07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    strcpy(updstsmg, "#");

    /* send update status message to status table manager */
    msglen = strlen(updstsmg);
    sockfd = connectUN(smsoc);

    if ( ( writemsg(sockfd, updstsmg, msglen) < msglen) {
        printf("\tExecute Agreement: status update to sm failed\n");
        printf("\t07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    close(sockfd);
} /* end else !smallest */

} /* end if initcommitphase */

removelst(dn);
free(tpbuf);
free(stbuf);
free(gvbuf);

} /* end Execute Agreement */

```

```

        strcpy(initcommsg, "failcommit");
    else
        strcpy(initcommsg, "joincommit");

    strcpy(initcommsg, "");
    strcpy(initcommsg, subj);
    strcpy(initcommsg, "#");

    /* separator */
    /* subject for token */
    /* end of msg */

    /* send initiae commit msg */
    msglen = strlen(initcommsg);
    comfid = connectUN(comsoc);

    if ( ( writemsg(comfid, initcommsg, msglen) < msglen) {
        printf("\tExecute Agreement: init commit send failed for %s\n", subj);
        printf("\t07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    /* ### wait for response back from commit processor */

    if ( ( msglen=readmsg(comfid,&msg, "#") < 0 ) {
        printf("\tExecute Agreement: fail on response from commit\n");
        printf("\t07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    msg[msglen] = NULL;

    /* don't check message as only possible is completion flag from commit */

    close(comfid);

} /* end then (smallest == TRUE) */

else { /* !smallest -> update status with appropriate pending status */

    strcpy(updstsmg, "updstatus\n");
    strcpy(updstsmg, subj);
    strcpy(updstsmg, "");
}

```



```

/*****
OriginatorFailed:

    return TRUE if:

        RelativeRank(tp_origin, myaddr) < RelativeRank(subject(token), myaddr)
        and token is a 'failagree' token

    OR

        rank(tp_origin) > rank(myaddr) and token is a "joinagree" token

    otherwise return FALSE

*****/
int OriginatorFailed(token, gv, myaddr, origin)
char *token, *gv, *myaddr, *origin;
{
    char *subject, *tmp_tkn;
    link *list;
    int failed, rr_o, rr_s, rank_o, myrank, tokentype;

    tmp_tkn = CALLOC(strlen(token) + 1, char);
    strcpy(tmp_tkn, token);
    list = str2list(tmp_tkn, "\n #");

    if (getfromlist(list, &subject, 2) == 0) {
        printf("\nOriginatorFailed: token parsing for subject failed.\n");
        printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1);
    }

    failed = FALSE;

    tokentype = GetTokenType(token);

    if (tokentype == JOINAGREE) {
        if ((rank_o = GetRank(gv, origin)) < 0) {
            printf("\nOriginatorFailed: tp originator not in gv.\n");
            printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
            exit(-1);
        }

        if ((myrank = GetRank(gv, myaddr)) < 0) {
            printf("\nOriginatorFailed: i'm not in gv.\n");
            printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
            exit(-1);
        }

        if (rank_o > myrank) {
            failed = TRUE;
        }

        /* end if joinagree */

        if (tokentype == FAILAGREE) {
            if ((rr_o = RelativeRank(gv, origin, myaddr)) < 0) {
                printf("\nOriginatorFailed: tp originator not in gv.\n");
                printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
                exit(-1);
            }

            rr_s = RelativeRank(gv, subject, myaddr);

            if (rr_o < rr_s)
                failed = TRUE;

            /* end Failagree */

            remove(list);
            free(tmp_tkn);

            return(failed);
        }

        /* end OriginatorFailed */
    }
}

```

COMMIT PROCESSOR

Commit Processor - Process Dependencies	132
Actions for Committing a Change	133
Actions for <i>Committing a Change</i>	133
CommitProcessor	134

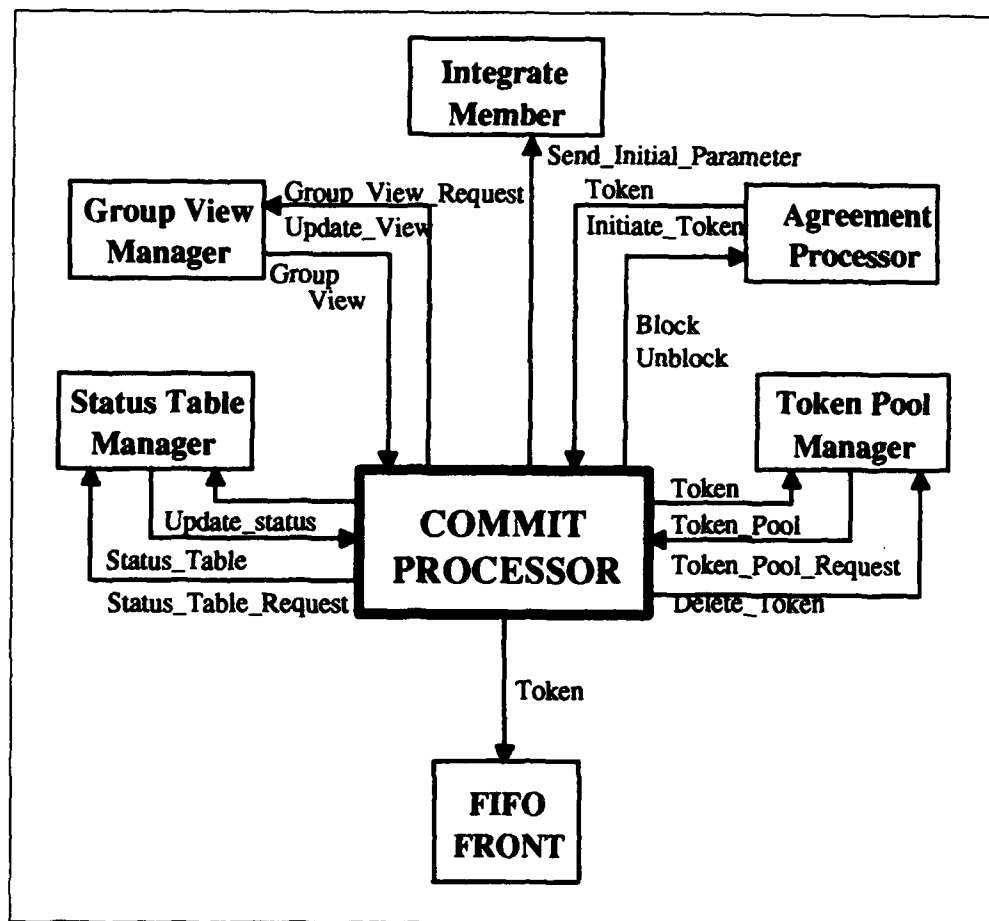


Figure A10 Commit Processor - Process Dependencies

```

CommitChange for  $commit_{p_j}(p_k)$  at  $p_i$ 
/* Depending on whether a join or departure */
1  add or delete  $p_k$  from  $GV(p_i)$ 
2  delete  $p_k$  entry from  $ST_{p_i}$ 
3   $vn(p_i) \leftarrow vn(p_i) + 1$ 
4  delete all  $commit$  tokens received before  $agree_{p_j}(p_k)$  from  $TokenPool(p_i)$ 
5  if (join committed &&  $joinreq_{p_j}(p_k) \in TokenPool(p_i)$  )
6    delete  $joinreq_{p_j}(p_k)$ 
7  end
8  add  $commit_{p_j}(p_k)$  to  $TokenPool(p_i)$ 
9  delete  $agree_{p_j}(p_k)$ 
10 if (current host =  $p_k$ )
11   determine new  $p_{host}$ 
12 end
13 if ((join committed) && ( $p_{host} = p_i$ ))
14   send  $ST_{p_i}$ ,  $TokenPool(p_i)$ , and  $GV(p_i)$  to  $acwnbr(p_i)$ 
15 end
16 send  $commit_{p_j}(p_k)$  token to  $cwnbr(p_i)$ 

end CommitChange

```

Figure A11 Actions for Committing a Change

```

ProcessCommitTkn for  $commit_{p_j}(p_k)$  at  $p_i$ 
1  if (initiate commit message received)
2    generate commit token
3    token to be processed  $\leftarrow$  generated token
4  else if ( $(p_i \neq p_j)$  && (not duplicate))
5    token to be processed  $\leftarrow$  received token
6  else
7    exit
8  end
9  CommitChange
10 while (  $p_i \in ST_{p_i}$  with pending status &  $Rank(p_i) < Rank(p_m)$ ,  $p_m \in ST_{p_i}$  )
11   generate commit token
12   token to be processed  $\leftarrow$  generated token
13   CommitChange in rank order
14 end

```

Figure A12 Generate / Receive and Process a Commit Token


```
/* parse status table - add all members to list in rank order */
sl_ptr = 3;
finish = FALSE;
```

```
while (finish != TRUE) {
```

```
if ( getfromlist(slist, &st_subj, st_ptr) != st_ptr)
```

133

```
/* get the status */
```

```
if (getfromlist(slist, &status, st_ptr+1) != (st_ptr + 1)) {
```

```
printf("vorder pending: status parsing failed\n");
```

```
print("v07ssssssssssssssssssssssn");
```

```
/* check for joined status ... ignore it */
```

```
if (strcmp(status, "joined") != 0) {
```

```
/* get the corresponding agree token */
```

done = FALSE;

MD DR = 3:

```
while (done != TRUE){
```

```
/* check for end of token pool */
```

```
if (getfromlist(tplist, &next token, tp ptr) != tp ptr) {
```

```
printf("vorder pending: token not found\n");
```

```
printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSS");
```

```
exit(-1); }
```

else {

```
/* check for proper token */
```

```
tktemp = CALLOC(strlen(next_token) + 1, char);
```

```
strcpy(iktemp, nxt_token);
```

```
tklist = str2list(tktemp, "\n #");
```

```

/* find insertion point */
last = head;
point = last->next;

while ( (point != NULL) && (rank > point->rank) ) {
    last = point;
    point = point->next; }

/* create commit token */
new_token = CALLOC(TOKENMSGLLEN+1, char);

if ((strcmp(status, "joinpendg") == 0) || (strcmp(status, "failpendg") == 0)) {
    if (strcmp(status, "joinpendg") == 0)
        strcpy(new_token, "joincommit");
    if (strcmp(status, "failpendg") == 0)
        strcpy(new_token, "failcommit");

    strcat(new_token, tp_subj);
    strcat(new_token, " ");
    strcat(new_token, myaddr);
    strcat(new_token, "#");
}
else {
    strcpy(new_token, "agree");
}

/* don't insert a token for the corresponding commit token */
if (strcmp(tp_subj, subject) != 0) {
    /* insert in order */
    temp = make_node(rank, new_token);
    temp->next = last->next;
    last->next = temp;
}
free(new_token);
removelst(tklst);
free(tltemp);
} /* end if status != joinrqstd */

```

```

/* increment pointer for next status */
st_ptr = st_ptr + 2;

} /* end if next status */
} /* end while !finish */

removelst(tkn_lst);
removelst(tplst);
removelst(slist);
free(temp_tk);
free(temp_tp);
free(temp_st);

/* all tokens w/ subj in status table now ordered
 * discard all past first "agree"
 *
 * NOTE: the first "agree" is NOT the corresponding agree for this
 * commit.
 */
last = head;
point = last->next;
finish = FALSE;

while (finish != TRUE) {
    if (point == NULL)
        finish = TRUE;
    else {
        if (strcmp(point->data, "agree") == 0) {
            /* remove the end of the list */
            last->next = NULL;
            while (point != NULL) {
                ptr = point;
                point = ptr->next;
                free(ptr->data);
                free(ptr);
            }
        }
    }
}

```



```

free(gview);
} /* end else assemble update view */

/* assemble update status (delete) message */
strcpy(updstst, "updstatus\n");
strcat(updst, subject);
strcat(updst, " delstatus ");

msglen = strlen(updst);
smfId = connectUN(smsoc);

if ( writemsg(smfid, updst, msglen) < msglen ) {
    printf("\nCommit_Change: update status tx failed\n");
    printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }
close(smfid);

/* get local token pool */
tpool = GetTokenPool(tpmsoc);

DeleteTokens(tpool, token, tpmsoc);

/* send token to TPM */
token = CALLOC(strlen(token) + 9 + 1 + 1, char);
strcpy(token, "token\n");
strcat(token, token);

msglen = strlen(token);
tmfid = connectUN(tpmsoc);

if ( writemsg(tmfid, token, msglen) < msglen ) {
    printf("\nCommit_Change: token tx failed\n");
    printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }
close(tmfid);

/* assemble delete agree token message */
strcpy(deltkn, "deltoken\n");

if ( strcmp(ctype, "failcomit") == 0 )
    strcat(deltkn, "failagree ");
else
    strcat(deltkn, "joinagree ");

strcat(deltkn, subject);
strcat(deltkn, " ");
strcat(deltkn, originator);
strcat(deltkn, "#");

msglen = strlen(deltkn);
tmfid = connectUN(tpmsoc);

if ( writemsg(tmfid, deltkn, msglen) < msglen ) {
    printf("\nCommit_Change:delete token tx failed\n");
    printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }
close(tmfid);

if ((joine) && (strcmp(host, myaddr) == 0)) {

    /* ensure that the changes have been incorporated in the
       token pool prior to sending the initial parameters
       NOTE: async consideration */

    free(tpool);
    tpool = GetTokenPool(tpmsoc);

    /* send SendInitialParameters */
    strcpy(srndinp, "srndinpar\n");
    strcat(srndinp, subject);
    strcat(srndinp, "#");

    msglen = strlen(srndinp);
    initmfrid = connectUN(initmbsoc);

```



```

if ( writemsg(nimbrfid, sndinip, msglen) < msglen ) {
    printf("\nCommit_Change: sndinip tx failed\n");
    printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

close(nimbrfid);

} /* end if srcmp */

free(host);

/* send token to FRONT -> cwnbr */
msglen = strlen(token);
frontid = connectUN(frontsoc);

if ( writemsg(frontid, token, msglen) < msglen ) {
    printf("\nCommit_Change: token tx to front failed\n");
    printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

close(frontid);

removelst(tklist);
free(tktemp);
free(tktoken);
free(tpool);

} /* end Commit_Change */

```

```

/*****
DeleteTokens:
*****/

void DeleteTokens(tpool, token, tpmsoc)
char *tpool, *token, *tpmsoc;

{
    int    tpsize, i, msglen, tpnrfid;
    char   *tpktemp, *poolsz, *tktemp, *subject, *nxtkn, *nxtktemp,
            *nxttype, *nxtsubj, *nxtorig, deltkn[DELTKNLEN+1];
    link   *tplist, *tklist, *nxtklist;

    /* disassemble token pool */
    tpktemp = CALLOC(strlen(tpool) + 1, char);
    strcpy(tpktemp, tpool);
    tplist = str2list(tpktemp, "\n#");

    if ( getfromlist(tplist, &poolsz, 2) != 2 ) {
        printf("\nProcCommitPending: token pool size parsing failed\n");
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    tpsize = atoi(poolsz);

    /* get token subject */
    tktemp = CALLOC(strlen(token) + 1, char);
    strcpy(tktemp, token);
    tklist = str2list(tktemp, "\n#");

    if ( getfromlist(tklist, &subject, 2) != 2 ) {
        printf("\nProcCommitPending: token subject parsing failed\n");
        printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }

    for (i=3; i<=tpsize + 2; i++) { /*for all tokens in tpool*/

```


INTEGRATE MEMBER PROCESS

Process Dependencies	145
Process Specification	145
Integrate Member Code	146




Figure A13 Integrate Member - Process Dependencies

IntegrateMember

```

1  if (initial parameters)
2    send blocking message to status reporter
3    send GV to group view manager
4    send unblocking message to status reporter
5    send ST to status table manager
6    send TokenPool to token pool manager
7  else
8    get  $GV_{p_i}$  from group view manager
9    get  $ST_{p_i}$  from status table manger
10   get  $TokenPool(p_i)$  from token pool manager
11   assemble initparam message
12   send initparam message to new member
13  end
  
```

Figure A14 Integrate Member Process Specification

```

/*****
* Integrate Member (INTMBR)
*
* Version: 19 May 1993
* Author: David Pezdirtz
*
* DESCRIPTION:
*   * Waits for a connection and then reads the initial parameter msg.
*   * Sends the init sub-msg to the 3 data base managers. On initial param
*   * request, compiles the init_params and returns the init_param msg.
*
* USAGE: numbr soc my_addr gvmnsoc smsgsoc tpmnsoc backsoc
*
*****/

/*****
Declarations
*****/
#include "gmp.h"
#include "msgutil.c"
#include "socutil.c"

typedef char *buffer;

void demux0;
void mux0;

/*****
Main:
*****/
main(argc,argv)
int  argc;
char *argv[];
{
    soc, newsoc, clen, msglen;
    *gvmnsoc, *smsgsoc, *tpmnsoc, *backsoc;

```

```

char      *my_addr, *target;
int
struct sockaddr_un
char
link
*msg11, *msg12;

if ( argc == 7 ) {
    soc = atoi(argv[1]);
    my_addr = argv[2];
    gvmnsoc = argv[3];
    smsgsoc = argv[4];
    tpmnsoc = argv[5];
    backsoc = argv[6]; }

else {
    printf("INTMBR usage error: numbr soc my_addr gvmnsoc smsgsoc tpmnsoc backsoc\n");
    printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
    exit(-1); }

while( TRUE ) {
    clen = sizeof(caller_addr);

    /* Accept connection request from client */
    clen = sizeof(caller_addr);

    if ( (newsoc = accept(soc, (struct sockaddr*) &caller_addr, &clen)) < 0 ) {
        printf("INTMBR: accept error\n");
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1); }

    /* Read message */
    if ( (msglen = readmsg(newsoc, &buf, "#")) < 0 ) {
        printf("INTMBR: read error\n");
        printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1); }

    buf[msglen] = NULL;

```

```

/* close the input socket */
close(newsoc);

/* Determines which message was received */
action = in_msg_type(buf);

/* set up the top level list */
msgl1 = str2list(buf, "\n@#");

switch (action) {
case INITPARAM:
    demux(msgl1, gvmsock, smsock, tpmsoc);
    break;

case SNDINIPAR:
    if (getfromlist(msgl1, &target, 2) != 2) {
        printf("INTMBR error: get target\n");
        printf("\07SXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
        exit(-1); }

    mux(backsoc, gvmsock, smsock, tpmsoc, target, my_addr);
    break;

default:
    printf("INTMBR error: invalid msg type\n");
    printf("\07SXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
    exit(-1);
} /* end switch */

/* Dispose of msg list */
removelst(msgl1);
free(buf);

} /* end while */

} /* end main */

```



```

if (writemsg(backfd, message, message_len) < message_len) {
    printf("Integrate Mbr unix: writemsg -- still no go\n");
    printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
    exit(-1); }

close(backfd);

/* clean up memory allocated */
removelst(gvmsgl);
removelst(smmsgl);
removelst(tpmsgl);

free(message);
free(gvmsgl);
free(smmsgl);
free(tpmsgl);

} /* end mux */

```

```

tpmsgl = str2list(tpmsg, "\n#");

if ( getfromlist(gvmsgl, &gv, 2) == 0) {
    printf("Integrate Mbr error: getfromlist gv\n");
    printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
    exit(-1); }

if ( getfromlist(smmsgl, &st, 2) == 0) {
    printf("Integrate Mbr error: getfromlist st\n");
    printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
    exit(-1); }

if ( getfromlist(tpmsgl, &tp, 2) == 0) {
    printf("Integrate Mbr error: getfromlist tp\n");
    printf("07XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX\n");
    exit(-1); }

/* reserve temp storage for outgoing message */
message = CALLOC((9 + 2*MAXLMTSIZE + strlen(gv) + strlen(st) + strlen(tp) + 5 + 11),
char);

/* assemble the message */
strcpy(message, "initparam\n");
strcat(message, target);
strcat(message, "\n");
strcat(message, my_addr);
strcat(message, "@\n");
strcat(message, gv);
strcat(message, "@\n");
strcat(message, st);
strcat(message, "@\n");
strcat(message, tp);
strcat(message, "#");

message_len = strlen(message);

/* write the message */
backfd = connectUN(backsoc);

```

JOIN PROCESSOR

Process Specification	152
Join Processor Code	153

InitiateJoin for a join request message/token for p_{new} at p_i

```

1  while (true)
2    if ( $p_{new} \in ST_{p_i}, GV_{p_i}$ )
3      receive join request message or token for  $p_{new}$ 
4    end
5    if ( $p_i = p_{host}$ )
6      send initiate agreement message to AgreeProcessor for  $p_{new}$ 
7      block until AgreeProcessor acknowledges end of processing
8    else
9       $ST_{p_i}(p_{new}) \leftarrow JoinRequested$ 
10     if (join request message) /*  $p_{new}$  locates  $p_i$  and sends its join request */
11       generate  $joinreq_{p_i}(P_{new})$  token
12     end
13     add  $joinreq_{p_i}(p_{new})$  to  $TokenPool(p_i)$ 
14     send  $joinreq$  token to  $cwnbr(p_i)$ 
15   end
16 end

end InitiateJoin

```

Figure A15 Processing of a Join Request Message / Token


```

targetmember[j] = NULL;

/* assemble a request */
strcpy(joinrequest, "joinrequest\n");
strcat(joinrequest, targetmember);
strcat(joinrequest, "\n");
strcat(joinrequest, myaddr);
strcat(joinrequest, "#");

/* send a join request */
msglen = strlen(joinrequest);
backfd = connectUN(backsoc);

if ( writemsg(backfd, joinrequest, msglen) < msglen) {
    printf("\nJOINPROCESSOR: join request to back failed.\n");
    printf("07$");
    exit(-1);
}
close(backfd);

/* Make soc nonblocking to enable count down while waiting */
on = NONBLOCKING;
ioctl(soc, FIONBIO, (char *) &on);
countval = GetCountForJoin();

while ( (foundmember == FALSE) && (CountDown(&.countval,
COUNT_DOWN_STEP) != 0) ) {
    clen = sizeof(caller_addr);

    if ( (newsoc = accept(soc, (struct sockaddr *) &caller_addr, &clen)) == -1) {
        if ( (errno == EWOULDBLOCK) || (errno == EINPROGRESS) ) {
            /* no connections present */
        } else {
            printf("\nJOINPROCESSOR: initparam accept error.\n");
            printf("07$");
            exit(-1);
        }
    } /* end if newsoc */
    else { /* a connection is present */

```

```

frontsoc = argv[9];
groupname = argv[10];
sites = argv[11];

sitelist = str2list(sites, "=");

if ( (nsites = listsize(sitelist)) == 0 ) { /* i must be >= 1 */
    printf("\nJOINPROCESSOR: site list parsing failed or empty list.\n");
    printf("07$");
    exit(-1);
}

i = 0;
foundmember = FALSE;

while ( (foundmember == FALSE, && (++i <= nsites) ) ) {
    if ( getfromlist(sitelist, &sitename, i) == 0 ) {
        printf("\nJOINPROCESSOR: sit.: list parsing failed for site %d.\n", i);
        printf("07$");
        exit(-1);
    }

    if ( CopyGVFile(sitename, myaddr, groupname) == 0 ) {
        printf("\nJOINPROCESSOR: group view file not found at %s.\n", sitename);
    } else { /* file found and copied, if found assumed to be in the correct format */
        fileptr = fopen(groupname, "r");
        fscanf(fileptr, "%d", &groupsize);
        fscanf(fileptr, "%d", &groupsize); /* overwrite */

        c = fgetc(fileptr); /* to get rid of the \n after groupsize */

        while ( (foundmember == FALSE) && (groupsize-- > 0) ) {
            /* get the next entry as the request target */
            j = 0;

            while ( ( (c = fgetc(fileptr)) != EOF) && (c != '\n') )
                targetmember[j++] = c;

```



```

*****
* Element has become a member of the name group. Now, simply wait for
* join requests or tokens to arrive and process them.
*****
while (TRUE) {

    /* read join request from front or join token from agreement process */
    clen = sizeof(caller_addr);

    if ((newsoc = accept(soc, (struct sockaddr*) &caller_addr, &clen)) < 0) {
        printf("JOINPROCESSOR: unix accept error for join req / token.\n");
        printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    if ((msglen = readmsg(newsoc, &buf, "#")) < 0) {
        printf("JOINPROCESSOR: unix read error for join request or token.\n");
        printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    buf[msglen] = NULL;
    close(newsoc);

    /* only tokens rcvd here are join tokens or joinrequests */
    msgtype = in_msg_type(buf);

    if ((msgtype != JOINREQST) && (msgtype != TOKENOKN)) {
        printf("JOINPROCESSOR: join req or token expected. rcvd => %s\n", buf);
        printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    /* get the originator of the request or subject of token
    * (3rd field in both
    */
    tmp_buf = CALLOC(strlen(buf) + 1, char);
    strcpy(tmp_buf, buf);

    msg = str2list(tmp_buf, "\n #");
}

if (getfromlist(msg, &originator, 3) == 0) {
    printf("JOINPROCESSOR: group join request/token parsing failed.\n");
    printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1);
}

/* if I'm the host -> initiate joinagree token */
if (InGroup(gvmsoc, myaddr) == 0)
    SendTk2Agr(agrsoc, "joinagree", originator);
else {
    /* add joinreqst status to status table */
    strcpy(updstismsg, "updstatus\n");
    strcat(updstismsg, originator);
    strcat(updstismsg, "\n");
    strcat(updstismsg, "joinpendg\n");
    strcat(updstismsg, "#");

    /* send update status message to status table manager */
    msglen = strlen(updstismsg);
    smfd = connectUN(smsoc);

    if (writemsg(smfd, updstismsg, msglen) < msglen) {
        printf("VExecuteAgreement: status update to sm failed.\n");
        printf("V07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1);
    }

    close(smfd);

    /* if joinreq msg -> generate joinreqst token */
    if (msgtype == JOINREQST) {
        token = CALLOC(TOKENMSGLEN + 1, char);
        strcpy(token, "tokenokn\njoinreqst");
        strcat(token, originator);
        strcat(token, "\n");
        strcat(token, myaddr);
        strcat(token, "\n");
        token = CALLOC(strlen(buf) + 1, char);
        strcpy(token, buf);
    }
}

```

```

/*****
CopyGVFile: Returns a 0 if file not found at the site ip address sitename.
Else returns a 1 after copying the file in the local directory
with name as groupname.
*****/
int CopyGVFile(sitename, myaddr, groupname)
char *sitename, *myaddr, *groupname;
{
    int copied;
    u_long inaddr;
    char command[512], *ip_address, localaddr[MAXLMTSIZE+1], *myname;
    FILE *fileptr;
    struct hostent hentry, *hptr;
    link *myaddrlist;

    hptr = &hentry;

    if ((inaddr = inet_addr(sitename)) != INADDR_NONE) {
        bcopy((char *) &inaddr, (char *)hptr->h_addr, sizeof(inaddr)); }
    else {
        if (hptr=gethostbyname(sitename))
            printf("\tCopyGVFile: Success, found %s, also known as %s\n",
                hptr->h_name, hptr->h_aliases[0]);
        else
            printf("\tCopyGVFile: Sorry host %s not found\n", sitename);
    } /* end else inaddr */

    ip_address = (char *) inet_ntoa( *((struct in_addr *) (hptr->h_addr)));

    strcpy(localaddr, myaddr);
    myaddrlist = str2list(localaddr, ",");

    if (getfromlist(myaddrlist, &myname, 1) != 1) {
        printf("\tCopyGVFile: my addr parsing failed\n");
        printf("\t07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
        exit(-1); }
}

```

```

/* add joinreqst token to token pool */
msglen = strlen(token);
tokenid = connectUN(qmsoc);

if ( writemsg(tokenid, token, msglen) < msglen) {
    printf("\tProcessToken: tkn send to join process failed\n");
    printf("\t07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

close(tokenid);

/* send joinreqst token to cwnbr */
msglen = strlen(token);
tokenid = connectUN(frontsoc);

if ( writemsg(tokenid, token, msglen) < msglen) {
    printf("\tAgreement: token send to front failed\n");
    printf("\t07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

close(tokenid);

free(token);

} /* end else InGroup */

removeid(msg);
free(buf);
free(tmp_buf);

} /* end while true */

} /* end JOINPROCESSOR */

```



```

if ( strcmp(myname, ip_address) != 0 ) { /* need to attempt copying */
    sprintf(command, "rsh %s cat %s > %s", sitename, groupname, groupname);

    if ( system(command) != 0 ) {
        printf("\nCopyGVFile: Could not rsh at site %s\n", sitename);
        copied = 0;
    }
    else {
        sprintf(command, "chmod 777 %s", groupname);

        if ( system(command) != 0 ) {
            printf("\nCopyGVFile: Could not change mode.\n");
            printf("07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
            exit(-1);
        }
    } /* end else system */
} /* end if strcmp */

fileptr = fopen(groupname, "r");

if (fileptr) {
    if (fgetc(fileptr) != EOF)
        copied = 1;
    else copied = 0;
    fclose(fileptr);
}
else copied = 0;

remove(list(myaddress));

return(copied);

} /* end CopyGVFile */

```

```

/*****
GetCountForJoin: Determines the wait before the join request is repeated.
(value returned times the COUNT_DOWN_STEP seconds)
*****/
int GetCountForJoin()
{
    int next;

    next = RESENDREQST;
    return(next);
} /* end GetCountForJoin */

```

DATABASE MANAGERS

Group View Manager Process Dependencies	160
Status Table Manager Process Dependencies	160
Token Pool Manager Process Dependencies	161
Group View Manager	162
Status Table Manager	169
Token Pool Manager	174




Figure A16 Group View Manager - Process Dependencies

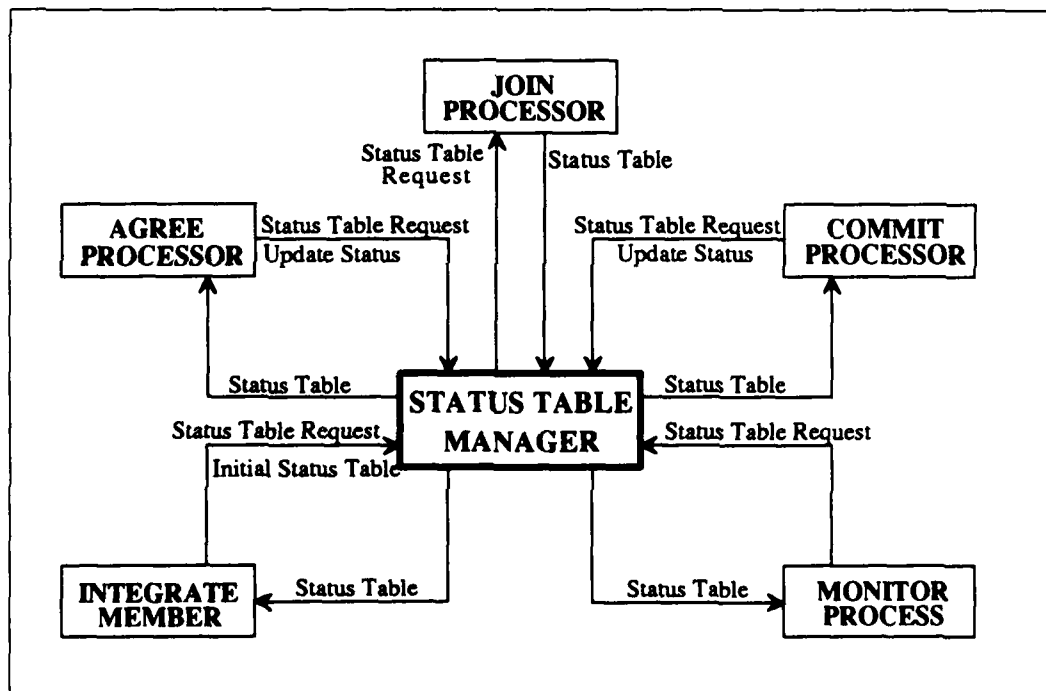


Figure A17 Status Table Manager - Process Dependencies




Figure A18 Token Pool Manager - Process Dependencies

```

/*****
 * Group View Manager (GVM)
 *
 * Version: 06 APR 1993
 * Author: David Pezdirtz
 *
 * DESCRIPTION:
 *   Waits for a connection and then reads a message. Depending on the
 *   message type it executes one particular action. Dumps the group view
 *   to a group file and sends the group view to the application socket on
 *   every update to the group view.
 *
 * USAGE: gvm soc init_element appsoc group_name
 *
 *****/

/*****
 * Declarations
 *****/

#include "grp.h"
#include "msgutil.c"
#include "socutil.c"

#define ADD      1      /* possible update actions */
#define DEL     2

typedef char *buffer;

void create_msg0;
void process_update0;
void process_request0;
void initialize0;
void add0;
void del0;
void remove_view0;
void save_grp_view0;

void show_grp_view0;

```

```

/*****
Main :
*****/

main(argc,argv)
int  argc;
char *argv[];
{
    int          soc,newsoc,clen,msglen,appid;
    int          action;
    struct sockaddr_un caller_addr;
    char          *msgtype,*buf,*tmp,*grp_name,*appsoc;
    link          *init_element;
    link          *head=NULL;
    link          *msgl;
    int          view_nmr=0;
    int          view_size=1;

    if (argc==5){
        soc = atoi(argv[1]);

        tmp = CALLOC(HEADERSIZE + 4 + strlen(argv[2]), char);
        strcpy(tmp,"inithead\n0=1=");
        strcat(tmp,argv[2]);

        appsoc = argv[3];
        grp_name = argv[4];
        else {
            printf("GVM usage error: gvm soc init_element app_soc grp_name\n");
            printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
            exit(-1);
        }

        /* initialize the group view for a single member */
        init_element = str2list(tmp,"n");
        initialize(&head, init_element, &view_nmr, &view_size);

        remove(list(init_element);
        free(tmp);

```



```

/*****
process_request : assemble the list into a string and write the message to
the socket

output msg format:
"groupview"\n | view # | = | view size | = | element | = | element | ... | # |

*****/
void process_request(soc, view, nmb, size)
int soc;
link *view;
int nmb, size;
{
    char temp[MAXNUM + 1];
    char *header, *msg;
    int msglen;

    header = CALLOC(9 + 2*MAXNUM + 3 + 1, char);
    msg = CALLOC((9 + 2*MAXNUM + 3 + size*(MAXLMTSIZE + 1) + 1), char);

    /* create the msg header */
    strcpy(header, "groupview");
    strcat(header, "\n");

    sprintf(temp, "%i", nmb);
    strcat(header, temp);

    strcat(header, "=");
    sprintf(temp, "%i", size);
    strcat(header, temp);

    /* create the message */
    msg = list2str(view, header, "=", "=");

    strcat(msg, "#");

    msglen = strlen(msg);

```

```

if ((writemsg(soc, msg, msglen)) != msglen) {
    printf("GVM process_request error: writemsg\n");
    printf("07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
    exit(-1);
}

free(msg);
free(header);

} /* end process_request */

/*****
add : add a node to the end of the current group view.

*****/
void add(head, element)
link **head;
char *element;
{
    link *tmp, *ptr;

    tmp = CALLOC(1, link);
    tmp->data = CALLOC(strlen(element) + 1, char);
    strcpy(tmp->data, element);
    tmp->next = NULL;

    if (*head == NULL) /* check for empty list */
        *head = tmp;
    else /* parse to end of list */
        ptr = *head;

        while (ptr->next != NULL) {
            ptr = ptr->next;
        }

        ptr->next = tmp;

    } /* end else head */

} /* end add */

```



```

/*****
process_request : assemble the list into a string and write the message to
the socket

output msg format:

"status" \n | table size | = table status | = element ... | # |

*****/
void process_request(soc, table, size)
int soc;
link *table;
int size;

{
    char temp(MAXNUM + 1);
    char *header, *msg;
    int msglen;

    /* reserve memory for header */
    header = CALLOC((9 + MAXNUM + 2 + 1), char);

    /* reserve memory for msg */
    msg = CALLOC((9 + MAXNUM + 2 + size*(MAXMTSIZE + 9 + 1) + 1), char);

    /* create the msg header */
    strcpy(header, "status");
    strcat(header, "\n");

    sprintf(temp, "%i", size);
    strcat(header, temp);

    /* create the message */
    msg = list2str(table, header, "=", "=");

    strcat(msg, "#");

    msglen = strlen(msg);

```

```

if ((write(soc, msg, msglen)) != msglen) {
    printf("STM process request error: write\n");
    printf("\07SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS\n");
    exit(-1); }

free(msg);
free(header);

} /* end process_request */

/*****
add : add a node to the end of the current group view.

*****/
void add(head, element)
link **head;
char *element;
{
    link *tmp, *ptr;

    tmp = CALLOC(1, link);
    tmp->data = CALLOC(strlen(element) + 1, char);
    strcpy(tmp->data, element);
    tmp->next = NULL;

    if (*head == NULL)
        /* check for empty list */
        *head = tmp;
    else {
        ptr = *head;
        while (ptr->next != NULL) {
            ptr = ptr->next; }
        ptr->next = tmp;
    } /* end else head */
} /* end add */

```



```

/*****
process_request : assemble the list into a string and write the message to the
socket

output msg format:

"tokenpool"nsize| = token type| sp|subj| sp|origint| sp| = | ... |#|

*****/
void process_request(soc, size)
int soc;
int size;
{
    char temp[MAXNUM + 1];
    char *header, *msg;
    int msglen;

    /* reserve memory for header */
    header = CALLOC(9 + MAXNUM + 2 + 1, char);

    /* reserve memory for msg */
    msg = CALLOC(9 + MAXNUM + 2 + size*(MAXMLTSIZE + 9 + 1) + 1, char);

    /* create the msg header */
    strcpy(header, "tokenpool");
    strcat(header, "\n");
    sprintf(temp, "%i", size);
    strcat(header, temp);

    /* create the message */
    msg = list2str(header, header, "=", "=");
    strcat(msg, "#");

    msglen = strlen(msg);

    if ((write(soc, msg, msglen)) != msglen) {
        printf("TPM process_request error: write msg");
        printf("\n");
        exit(-1);
    }
}

/*****
free(msg);
free(header);

) /* end process_request */

*****/
add : add a node to the end of the current table view.

*****/
void add(element)
char *element;
{
    link *tmp, *ptr;

    tmp = CALLOC(1, link);
    tmp->data = CALLOC(strlen(element) + 1, char);
    strcpy(tmp->data, element);
    tmp->next = NULL;

    if (head == NULL) { /* check for empty list */
        head = tmp;
        tail = tmp;
    }
    else {
        tail->next = tmp;
        tail = tmp;
    }
    /* end add */
}

```

```

*****
del : delete a node from the current token pool. Searches the token pool for
the token that matches the data element passed. Ignores the originator
field for comparison purposes.

*****
void del(element)
char *element;
{
    link *ptr, *last = NULL;
    int found = FALSE;
    char *data, *cur_data;

    ptr = head;

    /* strip off the originator */
    data = strip(element);

    /* search for the proper status */
    while ((found == FALSE) && (ptr != NULL)) {

        /* strip the originator field from the stored tokens */
        cur_data = strip(ptr->data);

        /* compare the desired data to the current node's data */
        if (strcmp(data, cur_data) == 0)
            found = TRUE;

        else {
            last = ptr;
            ptr = ptr->next;
        }

        free(cur_data);

    } /* end while found & ptr */

    free(data);

    /* ERROR: requested delete not in list */
    if (found == FALSE) {
        printf("TPM delete token: token %s\n", element);
        printf("\075\n");
        exit(-1);
    }

    /* check for delete of 1st element */
    if (last == NULL)
        head = ptr->next; /* if so, reset the head ptr */
    else
        last->next = ptr->next; /* otherwise remove node from list */

    /* check for delete of last element */
    if (ptr->next == NULL)
        tail = last;

    /* free up the storage used */
    free(ptr->data);
    free(ptr);

} /* end del */

```


DATA CRUNCHING PROGRAM

```

/*****
* Cruncher.c
*
* Description: gather all data generated by simpleapp, condense and dump
* the output. Requires a group to grow to max size and then shrink back to
* a single member. All failing members MUST be the acwnbr of the host.
* i.e. the last member of the group.
*
* Usage: crunch <max group size> <file listing members> [<output file>]
* NOTE: if <output file> is omitted, the output is sent to stdout
*
* Written by: david pezdirtz
*
* last revision: 26 Nov 1993
*
*****/

#include <stdio.h>
#include <sys/filio.h>
#include <sys/time.h>
#include <time.h>

#define CALLOC(n,type) (type *)calloc((unsigned) n,sizeof(type))

main(argc,argv)
int argc;
char *argv[];
{
    char *mbr_file, *output, *cmd, host[10], type[6], type2[5], temp[3],
    filename[10], subj[30], origin[30], grp_view[10];
    FILE *mbr_fd, *out_fd, *test_fd;
    int max_mbr, i, index, skip, total, et, sum;
    float avg;
    long sec, usec;

    cmd = CALLOC(128, char);

```

```

switch (argc){
case 3:
    max_mbr = atoi(argv[1]);
    mbr_file = argv[2];
    output = NULL;
    break;

case 4:
    max_mbr = atoi(argv[1]);
    mbr_file = argv[2];
    output = argv[3];
    break;

default:
    printf("CRUNCH usage error: crunch <max # of mbrs> mbrfile [outfile\n");
    printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
    exit(-1);
}

/* open the member file */
mbr_fd = fopen(mbr_file, "r");

if (mbr_fd == NULL) {
    printf("CRUNCH: attempt to open mbr file\n");
    printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
    exit(-1); }

/* open the output file (if necessary) */
if (output != NULL){
    out_fd = fopen(output, "a");
    if (out_fd == NULL){
        printf("CRUNCH: attempt to open output file\n");
        printf("\07$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n");
        exit(-1); }
    }
else {
    out_fd = stdout;
}

```


INITIAL DISTRIBUTION LIST

	Number of Copies
1. Defense Technical Information Center Cameron Station Alexandria, Virginia 22304-6145	2
2. Library, Code 52 Naval Postgraduate School Monterey, California 93943-5101	2
3. Chairman, Code EC Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, California 93943-5121	1
4. Professor Shridhar B. Shukla, Code EC/Sh Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, California 93943-5121	2
5. Professor Randy L. Borchardt, Code EC/Bt Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, California 93943-5121	2
6. LT David J. Pezdirtz, Jr. 250 Fellows Ave. West Jefferson, Ohio 43162	2